



Jet Propulsion Laboratory

California Institute of Technology

Final Report NASA NIAC Phase I Study

GRAVITY POPPERS: HOPPING PROBES FOR THE INTERIOR MAPPING OF SMALL SOLAR SYSTEM BODIES

Prepared for
Jason Derleth
Program Executive, NASA Innovative Advanced Concepts Program

Submitted March 14, 2021

Submitted by
Dr. Benjamin Hockman, PI
Jet Propulsion Laboratory
California Institute of Technology
4800 Oak Grove Dr., Pasadena, CA 91109
Phone: (626)639 5505
Email: benjamin.j.hockman@jpl.nasa.gov

Prof. Jay McMahon, Co-I
Aerospace Engineering Sciences
University of Colorado Boulder
Email: jay.mcmahon@colorado.edu

Prof. Daniel Scheeres, Co-I
Aerospace Engineering Sciences
University of Colorado Boulder
Email: scheeres@colorado.edu

Kalind Carpenter, Co-I
Jet Propulsion Laboratory
California Institute of Technology
Email: kalind.c.carpenter@jpl.nasa.gov

Executive Summary

This study investigated a novel mission architecture for high-resolution gravity science on small Solar System bodies. The “Gravity Poppers” mission concept (see Fig. 5) consists of a swarm of small, minimalistic probes that would be deployed from a mother spacecraft to the surface of a small body and which would periodically and randomly hop around the surface. By tracking the ballistic sub-orbital trajectories of this hopping swarm, a very high resolution gravity field can be measured, which can then be used to answer a number of scientific questions of interest regarding the interior mass distribution of the body.

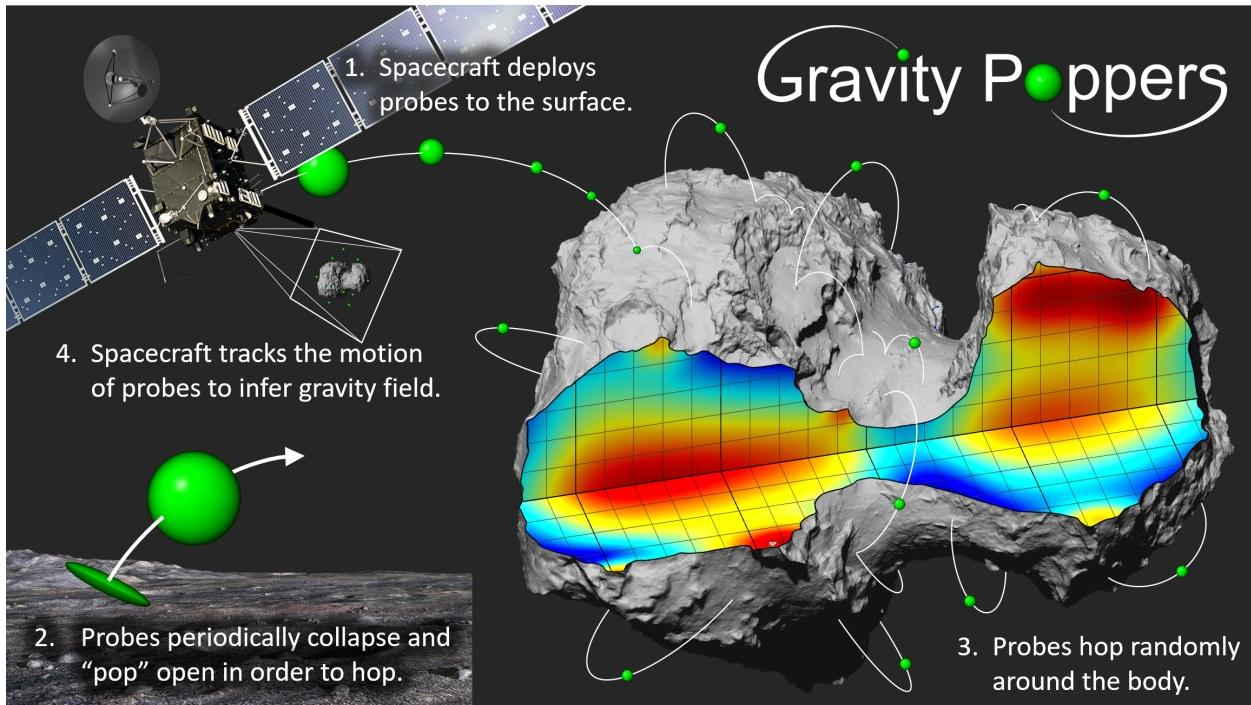


Figure 1: Gravity Popper mission concept: A spacecraft deploys a swarm of hopping probes to the surface of a small body and then tracks their trajectories to precisely and iteratively refine a model of the body’s gravity field and internal structure.

This concept was inspired, in part, by the recent serendipitous encounter of naturally-ejected particles from the surface of asteroid Bennu, which allowed the OSIRIS-REx team to estimate a much higher resolution gravity field than would have otherwise been possible. However, this process faced many challenges such as sparse tracking images and unknown particle properties, which ultimately limited its performance. The Gravity Poppers mission concept is an attempt to extend this fundamental concept of gravity estimation via tracked particles to the extreme by employing a swarm of artificial probes and tracking system that are specifically designed for this task. Thus, the Phase I feasibility study was not attempting to address whether or not enhanced gravity science is possible, but rather,

How accurately can the gravity field of a small body be estimated by tracking a swarm of hopping probes, and what are the key design parameters that most impact performance?

In other words, the Phase I study explored the highly coupled nature of gravity field ob-

servability and systems engineering aspects associated with the hopping probes and tracking system. Accordingly, the study had three main focus areas: (1) the probe tracking system, (2) the probe design and operation, and (3) swarm gravity estimation algorithms. A very large trade space of mission design parameters was revealed for each of these components (such as tracking modality, probe size, and swarm geometry). Key trades and sensitivity studies were performed for each of these components and a baseline point-design was evaluated for feasibility, performance, and cost. The results are extremely promising and summarized as follows:

Probe Tracking: The methods by which the trajectories of hopping probes can be measured was a key trade considered in Phase I, as it is the primary input to the gravity estimation filter. Various sensor architectures were considered, including both position and velocity sensors such as cameras, Lidar, and Doppler/ranging radios. Accordingly, three measurement modalities were considered: (1) angles-only from a camera on the mother spacecraft, (2) Doppler/ranging between the mother spacecraft and each probe, and (3) Doppler/ranging between all pairs of (local) probes. Different combinations of these sensors and modalities represents a trade-off of performance and complexity.

At first glance, it seemed as though an angles-only approach (i.e. imaging the swarm from a mother spacecraft) may be insufficient to measure the probes' trajectories, as there is an unobservable (depth) dimension. However, when leveraging the knowledge of dynamics through a batch estimation filter, an imaging-only approach performed surprisingly well, indicating that more complex ranging sensors may not be required. Increased imaging *frequency* (coupled with onboard image processing) was determined to be a key strength of this approach. A photometric analysis informed the requirements on relative *brightness* of the probes in order to be detected both against the dark sky as well as the lit body. Figure 2 shows an example of optical detection and tracking using a photo-realistic rendering of a bright probe in front of asteroid Bennu. Existing, small, astrometric-quality cameras were determined to be sufficient for optical tracking.

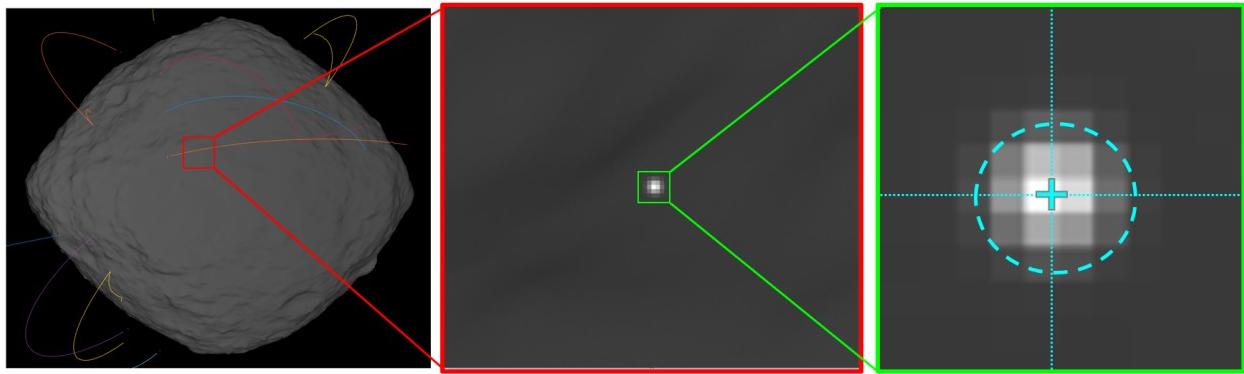


Figure 2: Simulation of optical tracking. **Left:** Probe tracks extracted by tracking algorithm on a synthetic (rendered) image of asteroid Bennu. **Center:** magnified synthetic image showing one probe's appearance in front of the small body's surface. **Right:** Visual representation of PSF centroiding.

Probe Design: The design of the hopping probes also had many design parameters to consider, which were driven by two fundamental requirements: that they (1) be “visible” to a camera on the mother spacecraft and (2) continuously hop around the body. Two probe

architectures emerged as promising candidates (see Fig. 3): (1) a white-painted sphere that reflects sunlight with a linear spring hop mechanism, and (2) a simple “femto-Sat” (roughly 5cm cube) with an array of strobed LEDs and an internal momentum actuator for hopping. Each design had different strengths that may be best suited to different mission scenarios, but both were sufficiently simple and small enough to accommodate 10’s of them in a small deployment canister.

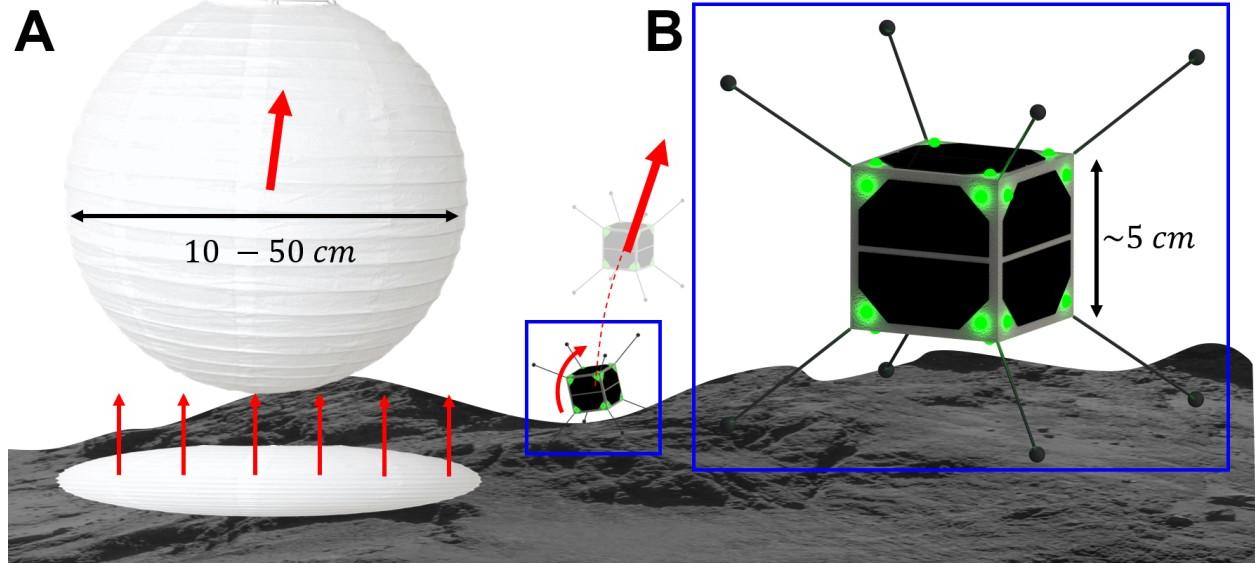


Figure 3: Two Gravity Popper probe design architectures derived in the phase I study. **(A)** A diffuse white sphere reflects sunlight to be detected by the spacecraft camera and uses an internal linear spring mechanism to hop. **(B)** A “femto-Sat” probe design which uses an omnidirectional array of strobing LED’s to be detected by the spacecraft camera and an internal momentum device (e.g. reaction wheel) for hopping.

Gravity Estimation: Finally, the *gravity estimation* was at the core of the Phase I study, in which we were able to leverage a state-of-the-art astrodynamics and orbit determination toolkit, MONTE, to demonstrate gravity recovery for realistic mission scenarios. In order to keep a tractable scope for Phase I, only a small subset of the many parameters that affect performance were explored through more detailed sensitivity studies. Ultimately, having explored only a small corner of this trade space (and with some conservative assumptions), we were able to demonstrate the ability to recover a ***degree-and-order 40*** spherical harmonic gravity field for a 500 m diameter body, corresponding to a mean spatial resolution of 20 m—sufficient to not only characterize a detailed map of the interior mass distribution, but also to “weigh” surface features such as boulders. Figure 4 shows the observability of a spherical harmonic gravity field from this study compared to OSIRIS-REx (both with and without particle tracking) as well as the spatial frequency of the corresponding density map that may be determined through gravity field inversion.

However, these results were constrained by computational limitations, and *extrapolation trends suggest that it would be possible to at least double the resolution* for longer simulations and with more computational resources.

Science Case Study: A given mass distribution produces a certain gravity field, but this field is not unique; that is, the same gravity field could be produced by other mass distri-

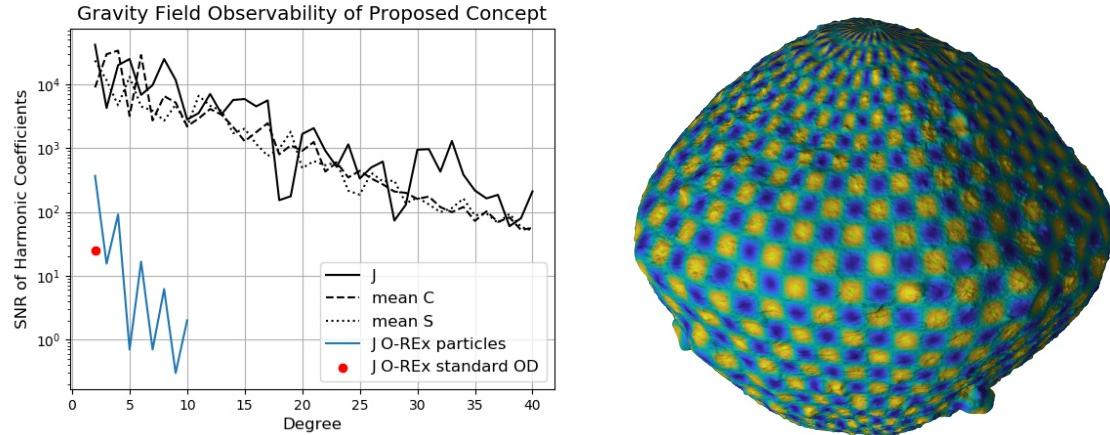


Figure 4: **Left:** Spherical harmonic SNR for a 24-hr simulation of 10 probes at a 10-s measurement frequency (black lines). OSIRIS-REx performance (with and without the use of particle tracking) is overlaid for reference. **Right:** The spatial frequency corresponding to a degree-and-order 40 gravity field, overlaid onto a model of asteroid Bennu, showing a minimum “wavelength” of about 20 m.

butions. Thus, the inversion of gravity fields to infer mass distribution is a famously under-determined problem that can yield many solutions. However, by leveraging observations of the surface morphology, geophysical constraints on interior structure, and potentially other types of measurements (e.g. radar and seismic), many scientific hypotheses can be tested with a gravity field of sufficient resolution. One such scientific question of interest that can be well constrained by a surface topographical shape model is the *density* of large boulders on the surface. Mass measurements of surface boulders may help to constrain the (perhaps heterogeneous) mineralogy of the body and provide further constraints to reason about the bulk composition and porosity of the interior.

We conducted a “boulder-weighing” experiment for a simulated mission to asteroid Bennu, in which the magnitude of a surface point-mass anomaly (or “mascon”), co-located with the center of a large boulder, was co-estimated with the global spherical harmonic gravity field. For a 24-hr simulation of 10 hopping probes and a 10-second measurement frequency, we were able to estimate the mascon μ parameter with high precision ($\text{SNR} > 10^3$) suggesting the ability to “weigh” large surface boulders with an error less than 0.1%.

In summary, The Phase I study revealed the immense potential of the Gravity Poppers mission concept to measure the gravity field of small bodies with exquisite precision, allowing for detailed inferences of the global mass distribution (two to three orders of magnitude better than existing methods). Such detailed gravity fields could help to address outstanding questions relating to origins science, planetary defense, and resource utilization. A systems engineering study of the hopping probes and tracking system also suggests that the hardware requirements are minimal and could be realized as a small, low-cost “instrument,” which could be hosted on a number of spacecraft visiting small bodies, thus having high potential for infusion. However, this Phase I study just scratched the surface, and many questions remain as to possible performance and potential applications of this concept. Future studies should (1) further explore the gravity field observability in more detail and with fewer simplifying

assumptions (e.g. for bodies of arbitrary shape), (2) increase the design maturity of small hopping probes, (3) demonstrate tracking performance with more realistic sensor models and hardware, and (4) apply this concept to a broader range of potential science applications.

Contents

1	Introduction	9
1.1	Background and Motivation	9
1.2	Current state-of-the-art for interior mapping	10
2	Mission Concept	11
3	Swarm Gravity Estimation	12
3.1	Objectives	12
3.2	Approach	13
3.3	Simulation Setup	13
3.3.1	Small Body	13
3.3.2	Mother Spacecraft	14
3.3.3	Probes	15
3.3.4	Landmarks	16
3.3.5	Optical Measurements	16
3.3.6	Dynamics Model	16
3.3.7	Filter Setup	17
3.4	Gravimetry Analysis	17
3.4.1	Baseline Covariance Analysis	18
3.4.2	Sensitivity to Imaging Frequency	19
3.4.3	Sensitivity to Probe Centroiding Precision	20
3.4.4	Ranging Measurements	21
3.4.5	Boulder Weighing Case Study	22
3.5	Numerical Considerations	23
4	Tracking System	24
4.1	Tracking Sensors Trade Study	24
4.2	Optical Tracking	25
4.2.1	Camera Selection	25
4.2.2	Spacecraft Distance	26
4.2.3	Probe Detection	26
4.2.4	Multi Object Tracking	28
5	Probe Design	30
5.1	Probe Requirements	31
5.2	Hopping System	32
5.2.1	Hopping Speed	32
5.2.2	Hopping Mechanism	35
5.2.3	Hopping Operations	36
5.3	Probe Visibility	37
5.3.1	Light Reflection	37
5.3.2	Light Emission	40
5.4	Power System	41

5.4.1	Hopping Power	41
5.4.2	Strobing Power	42
5.4.3	Hotel Power	43
5.4.4	Power System Architecture	44
5.5	Thermal System	44
5.6	Notional Probe Concepts	46
5.7	Probe Deployment System	48
6	Conclusions	48
7	Acknowledgements	50
8	References	51

1 Introduction

Project Objective: Develop a robust, low-cost mission architecture that enables the gravimetric density reconstruction of small body interiors (< 20 km) at spatial resolutions of 10–100 m, two to three orders of magnitude beyond existing methods.

1.1 Background and Motivation

The exploration of small solar system bodies has emerged as an important aspect of planetary exploration. Over the last two decades, the breadth and cadence of small body missions have grown considerably, including NEAR Shoemaker, Stardust, Hayabusa, Rosetta, Deep Impact, New Horizons, Dawn, Hayabusa2, OSIRIS-Rex, the most recent Discovery selections, Lucy and Psyche, and the SIMPLEx selection of Janus. The breadth of these missions reflects the tremendous diversity among “small bodies” (including asteroids, comets, and irregular moons), their broader interests to the scientific community for understanding the early formation and evolution of the solar system [1, 2], their potential for resource utilization [3], and the existential threat they pose to life on Earth [4]. However, despite the tremendous insight we’ve gathered from these missions (primarily through remote observations), many unanswered questions remain, particularly regarding their *interior composition and structure*.

As geophysical objects, understanding the interior structure of small bodies is an important goal for advancing our scientific knowledge of these objects. The presence of density variations due to mass concentrations and/or voids can provide important information and constraints on theories of how these small bodies formed and have evolved. Density inhomogeneities may also imply a varying composition – i.e. a rubble pile mixed of pieces of different mineralogy. Furthermore if we want to interact with these bodies more fully in the future for planetary defense or mining purpose, this knowledge will play a crucial role in conducting these operations. For example, different internal structures will influence the all-important momentum enhancement factor (β) which largely influences kinetic deflection missions, and mining operations will plan their operations with this information to determine their targets and avoid possible issues with voids.

Scientific studies have largely relied on geophysical models to make inferences about internal structure from sparse and imprecise data, such as how a body’s bulk density may indicate its “macroporosity” and thus, constrain its likely internal morphology (e.g. fractured solid, rubble pile, or dense core). Other studies have inspected meteorites to infer properties of near-Earth asteroids, despite their corrupted state from atmospheric entry [1, 2, 5, 6]. However, there is currently a severe lack of direct measurements to validate many of assumptions that underpin these models [2]. The information we have gleaned from missions – in particular NEAR, Hayabusa, Hayabusa2, OSIRIS-REx and Rosetta – have greatly informed our understanding, but is still limited to at best being able to infer slight deviations from homogeneous bodies [7]. We know from these results that there are important details and intricacies we have not yet been able to measure.

Unfortunately, measuring or inferring the interior structure of a small body is a particularly difficult task due to the small size and rubble pile nature. Currently, our most successful technique for gaining insights about the internal structure is through measurement/estimation of the body’s gravity field, which is a direct result of the mass distribution

and can be used to infer and constrain the internal structure, as discussed in the following section.

1.2 Current state-of-the-art for interior mapping

Historically, there are three primary geophysical methods that have been used to study the interior structure of Earth, and to some extent, other planetary bodies. However, they all pose significant challenges for studying the interiors of *small* bodies.

1. **Radar tomography** measures the transmission or reflection of radio signals from a surface and is used on Earth to image sinkholes, pipes, burial grounds, etc. Rosetta and Philae's CONCERT instruments were able to measure the dielectric constant of the small lobe of Comet 67P and infer bulk porosity [8], and some theoretical foundations have been developed for more detailed radar imaging of small bodies [9, 10]. However, the extension of this technique to a global imaging architecture requires instruments to be deployed from orbit, significantly reducing resolution and posing challenges with regard to data inversion and echo noise [1]. Also, certain asteroid materials such as clays and rocks flecked with metal may be opaque to electromagnetic energy [1].
2. **Seismic imaging** has taught us a great deal of the detailed structure of the Earth, but faces several challenges when translating this technology to small bodies. On the one hand, remote observations of the shift in regolith after a known impact/detonation event can help to infer crude estimates of the body's internal strength and local surface properties, as was done on the deep impact mission [11], Hayabusa2 [12], and the upcoming DART and Hera missions [13]. On the other hand, more refined tomographic imaging could be achieved through an array of surface seismometers and detonators, as proposed in a few mission concepts [14, 15, 16, 17]. However, a number of practical challenges remain for this concept including the precise deployment and localization of each surface probe, quantifying the true impulse imparted to the body from a surface charge, and chiefly, the mechanical coupling of seismometers to regolith-covered surfaces in microgravity.
3. Finally, **gravity measurements** offer an indirect way of probing a body's interior mass distribution, which has a long history of success on larger bodies such as planets and moons. Notably, through precise relative navigation, GRACE and GRAIL missions have produced gravity maps with unprecedented resolution on the Earth and Moon [18, 19]. However, the resolution of gravity models deduced from orbit determination (OD) is largely related to the size of the body and proximity to the surface. Dawn has been able to resolve harmonic gravity models for asteroids Vesta (500 km) and Ceres (1000 km) up to about degree and order 20 [20, 21], and NEAR was able to determine that the gravity model of asteroid Eros (17 km) agrees with a constant density assumption up to degree and order six [22]. However, gravity signals become very weak for bodies less than about 10 km, especially considering these bodies are often highly non-spherical, making it risky to orbit at low altitudes. For example, Rosetta and Hayabusa could only resolve a possible dual-density bi-lobed structure for comet 67P [23] and asteroid Itokawa [24], though even those results are controversial [2]. OSIRIS-REx was only expected to be able to measure up to fourth degree and order [25] using the spacecraft tracking data, however it ended up achieving a more precise gravity field estimate up to 10th degree due to the

serendipitous data acquired from observing ejected particles in Bennu’s vicinity [26, 7].

2 Mission Concept

Instead of adapting traditional planetary measurement approaches to small bodies, we propose a mission concept that is uniquely suited to the weak and irregular gravity fields of small bodies.

This proposed mission concept consists of a spacecraft which first rendezvous with a small body and enters into a distant orbit for a preliminary period of mapping and shape modeling—a common operation in small body missions. The spacecraft then lowers its orbit to an altitude at which it can deploy many small, minimalistic probes to the surface. The probes are periodically excited to produce a “momentum kick,” which induces large, random hops when resting on the surface. By precisely tracking the trajectories of the ballistic probes, the spacecraft can extract distributed gravity information and continuously refine the internal density model of the body. In other words, the probes need not be equipped with any sensors of their own; rather, they themselves act as the sensors through their observed ballistic response to the body’s gravity field.

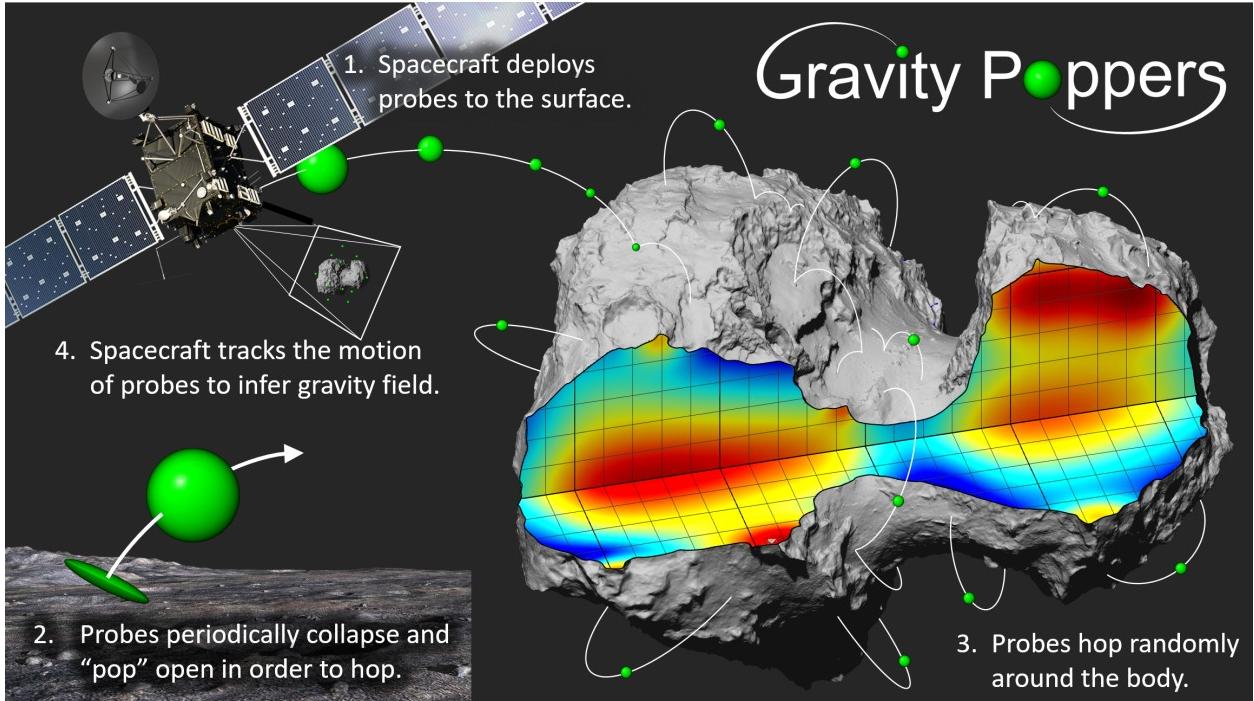


Figure 5: Gravity Popper mission concept: A spacecraft deploys a swarm of hopping probes to the surface of a small body and then tracks their trajectories to precisely and iteratively refine a model of the body’s gravity field and internal structure.

The key insight behind this approach is recognizing that the observability of subsurface density variations on small bodies can be dramatically improved by measurements that are (1) low altitude, (2) long duration, and (3) highly distributed around the body. Indeed, some of the most accurate gravity measurements of small bodies to date have been acquired during close proximity station keeping or touch-and-go sampling maneuvers [22, 24]. However,

due to the inherent risk of a primary spacecraft descending so close to the surface, these measurements are generally sparse. Thus, the key benefit of the *Gravity Poppers* mission concept is that it enables these close-proximity measurements with vastly greater spatial coverage through a robust, fault-tolerant swarm architecture. Furthermore, this mission concept has the strength of *extended time*: (1) a single ballistic trajectory may take hours to complete, allowing for more accurate tracking, and (2) since each hopper is long-lived (potentially executing hundreds of hops), the collective swarm yields a massive data set for more accurate density estimation.

The basic principle of gravity estimation via tracked suborbital particles was actually demonstrated through OSIRIS-REx's recent serendipitous encounter with the (cm-scale) rocks that were naturally ejected from the surface of Asteroid Bennu [26]. By tracking the trajectories of several particles through sparse imagery, a gravity field could be resolved up to degree 9—significantly higher than the degree-3 field that could be resolved with spacecraft gravity science alone. This gravity field could then be inverted to hypothesize a heterogeneous density distribution which could be explained by a low-density (porous) equatorial bulge [7]. However, impressive as these results were, this estimation faced many challenges that limited its performance such as the unknown size, shape, and optical properties of the particles (which had to be co-estimated), sparse imagery, and only few long-lived particle trajectory observations. Thus, it is expected that probes which are specifically designed to be tracked, with precisely calibrated mass, shape and photometric properties could be tracked far more precisely and over longer periods of time, resulting in a much more accurate gravity field estimation. The Gravity Poppers mission concept therefore offers a low-cost method for achieving dramatically enhanced gravity science at a broader range of small bodies which need not have active particle ejections of their own.

This Phase I project had three research pillars in order to address the key feasibility and maturation aspects associated with this concept: (1) swarm gravity estimation, (2) the tracking system, and (3) the design of the hopping probes. The remainder of this report is organized accordingly and concludes with a summary of the lessons learned and an outlook on the future studies to mature the concept for a flight mission.

3 Swarm Gravity Estimation

3.1 Objectives

The core of this study is to quantify gravity estimation performance for the Gravity Poppers mission concept. This mission architecture is characterized by a very large design trade space, such as number of probes, deployment strategy, hopping Delta-V, and probe tracking accuracy. These and other parameters directly affect the observability of the gravity field. Defining a mission design which optimizes gravimetry performance requires the development of high-dimensional sensitivity analyses and optimization of the trade space, which we are planning to tackle in future work. In Phase I, we evaluated a limited set of mission-design cases, mostly based on conservative assumptions, and assessed overall capabilities and limitations. In particular, the objectives of this study are the following:

1. Define one *baseline* mission design and OD setup, whose key parameters are in line

with previous, similar missions.

2. Identify the most influential parameters, in the OD setup and mission design, for gravity estimation.
3. Assess the potential of this mission concept for planetary science applications.

3.2 Approach

Our mission scenario is based on the OSIRIS-REx mission [27], which extensively studied the asteroid Bennu and collected a sample from its surface. Choosing Bennu as a central body in the simulations has two key advantages. First, this asteroid has been thoroughly characterized during said mission, and hence we use its physical properties such as mass, attitude dynamics, and size to simulate a ground truth which can be compared to OSIRIS-REx data. Second, at only 500 m in diameter, Bennu is a very small body and, as such, its gravity field is challenging to estimate as its gravitational signal is low relative to non-gravitational perturbations. From a gravity-signal standpoint, Bennu can be seen as part of the worst-case scenarios, within the small body spectrum, as opposed to more massive bodies (e.g., Comet 67P). Hence, we assume Bennu and its physical properties as the central body in our *baseline* simulation.

The remaining parameters in the OD setup are chosen following two key principles: (1) we choose parameters similar to the OSIRIS-REx mission when such values can be applicable to the Gravity Poppers case study, and (2) we make conservative assumptions for those parameters that are novel from this mission concept. Finally, the OD setup is informed by the mission design consideration for the onboard camera and probe design, discussed in Sections 4 and 5. The following sections describe in detail the key considerations for the OD setup.

To evaluate gravimetry performance while maintaining computational tractability, we firstly study a baseline scenario which resembles a realistic case, using conservative assumptions. Then, we compare this baseline with multiple cases where some design parameters are varied, to identify those which are most influential on overall gravity estimation. Such design considerations can be used to inform future case studies, where more thorough analyses to optimize the design will be carried out.

3.3 Simulation Setup

We implemented the dynamics and orbit determination simulations using the Mission Analysis, Operations, and Navigation Toolkit Environment (MONTE) [28], developed by NASA Jet Propulsion Laboratory. This software toolkit provides state-of-the-art astrodynamics and estimation libraries which are used for flight projects. The implemented software is divided into two subgroups: ground-truth simulation and orbit determination, which are executed in this order in the pipeline. The building blocks of the simulation pipeline are described in the following.

3.3.1 Small Body

The reference small body is asteroid 101955 Bennu, a near-Earth C-type asteroid. We assume Bennu's pole orientation of $[RA, Dec] = [85.65, -60.17]$ deg (in EME2000), rotation period

of $T = 4.296$ hours, and prime meridian PM = 54.77 deg. We also assume that the body is in principal axis rotation, i.e. no precession or nutation effects are accounted for. This is a good assumption for Bennu, which does not show any significant pole wobble [29]. The body's orbit is prescribed by an available SPK file on the NAIF SPICE repository and is assumed to be known a priori.

One simplification that we made in this Phase I study is to treat the *shape* of the body as an equivalent sphere. This was done for two reasons; first, the spherical harmonic representation of the gravity field we are estimating is known to have convergence issues within the circumscribing or “Brillouin” sphere of a body, and second, the collisional dynamics with an irregular polyhedron shape model was not readily available in the MONTE simulation environment we used for this study. Thus, to keep the analysis tractable for this Phase I study, we treated the shape of Bennu as sphere with radius of 246 m (Bennu's mean radius). However, we believe that while this simplifies the analysis, it is likely a conservative assumption for assessing gravity observability, as it does not allow probes to move within low-altitude concavities, which are much more gravitationally perturbed regions. Future work will need to consider other gravity field representations that converge within the Brillouin sphere (e.g., the interior Bessel gravity model [30]) and generalize the simulations to arbitrary shapes.

Bennu's point-mass parameter is $\mu = 4.892 \text{ m}^3/\text{s}^2$, whereas its non-uniform gravity is modeled using a standard spherical harmonic expansion. To generate a ground truth gravity model whose degree and order exceed the estimates from the OSIRIS-REx, we procedurally generate the normalized spherical harmonic coefficients. These are randomly sampled from the Kaula's power law for Bennu [25]:

$$C_{n0} = \frac{K_{zonal}}{n^{\alpha_{zonal}}}, \quad C_{rms} = \sqrt{\frac{\sum_{m=1}^n (C_{nm}^2 + S_{nm}^2)}{2n}} = \frac{K_{rms}}{n^{\alpha_{rms}}} \quad (1)$$

where C_{n0} are the normalized zonal coefficients (J) with degree n and C_{rms} are the normalized C and S coefficients. The values are randomly sampled from a normal distribution, where $K_{zonal} = 0.084 \pm 0.021$, $K_{rms} = 0.026 \pm 0.005$, $\alpha_{zonal} = 2.08 \pm 0.13$, and $\alpha_{rms} = 2.01 \pm 0.12$. The sign of the coefficient is also randomly sampled. Note that this gravity field does not coincide with the one estimated from Bennu (which was only characterized up to degree and order 8 [26], but uses Bennu's physical properties to define a power decay of its harmonics amplitude. Finally, we truncate the spherical harmonics at degree 40, to enable computational tractability. Future work will evaluate gravimetry performance for higher degrees and orders. The resulting ground truth spherical harmonics are shown in Figure 6.

3.3.2 Mother Spacecraft

The choice of observation geometry of the mother spacecraft is a complex trade space that is left for future work, and indeed may not even be under control of this system if it were to be hosted on a Spacecraft with other objectives. As discussed in Section 4, there is some flexibility to the observation distance of the spacecraft with an appropriate selection of optics (e.g. so that the small body fills some nominal fraction of the field-of-view). In this simulation study, we assume that the mother spacecraft is hovering at a distance of 4 km from the body's center of mass, to ensure good tracking performance with the baseline optical tracking system. The spacecraft trajectory is also *not* integrated (i.e. not dynamically

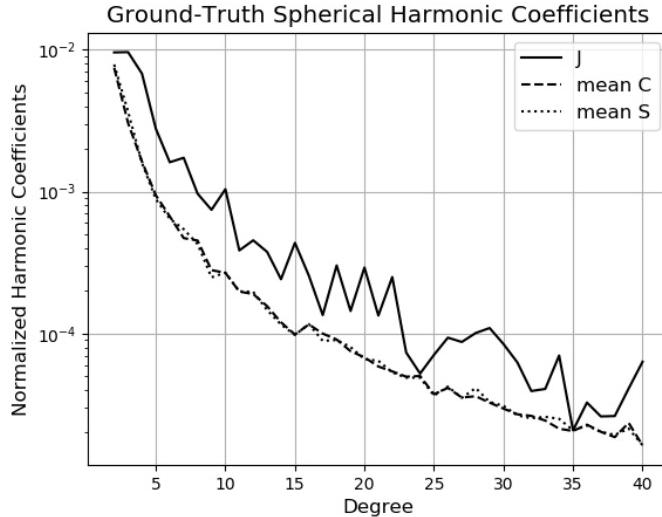


Figure 6: Amplitude of the normalized spherical harmonics coefficients, procedurally generated using random samples from Kaula’s power law for Bennu.

propagated). Instead we assume a constant position offset between the spacecraft and the small body, in the inertial frame, and treat the uncertainty of the spacecraft motion as a “consider” covariance, as discussed in Section 3.3.7 (Table 1). The hovering assumption is made to ensure an optimal viewpoint for probe tracking; future work will include trajectory design considerations for the mother spacecraft.

The spacecraft attitude model is nadir-pointing, with a stochastic pointing error of 0.003 degrees. This value corresponds to the one obtained during the OSIRIS-REx mission, after camera calibration. Hence, we assume good camera calibration in the following.

3.3.3 Probes

The specific designs of the hopping probes is discussed in Section 5. For this gravity estimation study, we need only define their mass and surface area. Here, we assume that the probes are spheres with a mass of 0.1 kg and radius 10 cm—corresponding to a worst-case scenario in terms of gravity observability due to their low ballistic coefficient, and hence a higher sensitivity to non-gravitational perturbation forces such as the solar radiation pressure (SRP) and albedo pressure. We assume a total of 10 probes, which is a highly notional estimate of the *minimum* number that may be accommodated in a realistic scenario. Clearly, the higher the number of probes, the higher is the expected surface coverage, for a given observation time frame.

The probes are deployed from the north and the south poles, using a random choice between the two. This deployment strategy maximizes the spatial coverage of the swarm, as there is a natural tendency for random hopping probes to migrate towards the lower geopotential latitudes. The probes’ motions are defined by an indefinite series of random hops, in which their velocities are randomly reinitialized upon collision with the surface. The hopping (exit) velocity relative to the surface is randomized both in magnitude and direction. The velocity magnitude is sampled from a normal distribution with mean $\mu = 0.3 \cdot v_{esc}$, and standard deviation $\sigma = 0.15 \cdot v_{esc}$ where $v_{esc} \approx 20 \text{ cm/s}$ is Bennu’s escape velocity at the

mean radius' altitude. This normal distribution is truncated to avoid very small and very big hops which are not physically realizable. The randomization of the probe velocity is meant to model variations due to the hopping mechanism precision and, chiefly, due to the interaction between the probe and the uncertain small body surface. The direction of the hop velocity is also randomized, sampling from a uniform distribution spanning between ± 30 degrees of the local vertical, to account for the rugged topography encountered by the probe.

3.3.4 Landmarks

In order to provide a realistic performance model of vision-based navigation, observations of surface landmarks are prescribed to provide body-fixed optical measurement to the navigation filter. Ten landmarks are generated by randomly sampling latitudes and longitudes. These are modeled as constant-offset trajectories with respect to the body center, defined in the body-fixed frame. Hence, they represent bias parameters to be estimated in the filter. Such landmarks could either be provided by natural terrain features on the surface using feature-detection algorithms or by the probes themselves that are resting on the surface.

3.3.5 Optical Measurements

Optical measurements are simulated using the ASTERIA camera specifications [31], an astrometric-quality CubeSat camera. Stellar aberration and geometric distortion are not considered here (and can be calibrated out in practice), and the exposure time is considered negligible. The imaging frequency for probe tracking is 10 seconds, which provides frame-to-frame measurement variations on the order of a few pixels. Simulated optical measurements are generated using the simulated trajectories of the probes and the camera, as well as using the body's shape to account for occlusions. Synthetic measurements are then perturbed with simulated white noise with a standard deviation of 0.05 pixels, i.e., about 5 times the expected centroiding error derived for a notional optical system in Section 4. Landmark measurements are assumed to have the same uncertainty, but are sampled with an interval of 5 minutes, in order to limit the total number of measurements to be processed by the filter.

3.3.6 Dynamics Model

The dynamics model used to propagate probe trajectories accounts for gravity, solar radiation pressure, and stochastic accelerations.

The gravity model includes both the attraction of the main solar system bodies and the spherical harmonic gravity of the small body. The solar system bodies considered are the Sun, the planets, and Pluto, modeled as point masses. The planetary ephemeris dataset DE430, from JPL's NAIF SPICE, is used to model the trajectory of such bodies. The spherical gravity of the small body is modeled by algorithmically generating spherical harmonic coefficients, as previously discussed.

The SRP is generated using a solar flux $\Phi = 101.98 \text{ kg km/s}^2$ and assuming a spherical shape of the probes. Future work will also include Thermal Radiation Pressure (TRP) and the Albedo pressure due to reflected photons from the small body's surface.

Stochastic accelerations are included to capture unmodeled forces acting on the probe dynamics, and are defined as zero-mean accelerations with a standard deviation of 10^{-12} km/s^2 —a value chosen based on the magnitude of the empirical (unmodeled) accelerations observed during the OSIRIS-REx mission. The accelerations batch size is chosen to be 0.1 hour, which

Name	Type	A-Priori σ
Pole RA	Bias	10 deg
Pole Dec	Bias	10 deg
Pole PM	Bias	10 deg
Pole ω	Bias	1 deg/day
SH coeff.	Bias	Given by Kaula's rule
Landmarks	Bias	0.1 m
Probes pos.	Dynamic	100 km
Probes vel.	Dynamic	100 m/s
S/C pos.	Consider	1 m
Stoch. Accel.	Stochastic	$1 \cdot 10^{-12}$ km/s ²
Stoch. Pointing	Stochastic	$1 \cdot 10^{-3}$ deg
Optical Msr.	N/A	0.05 pixel

Table 1: Navigation Filter Parameters

is one tenth of the one chosen for the OSIRIS-REx case [26]. We chose a smaller batch size as the sub-orbital arcs of the probes have a faster dynamics compared to the OSIRIS-REx observed particles. It should be noted that the batch size directly affects the OD results; future work will assess the optimal batch size for this mission concept.

3.3.7 Filter Setup

We use a square-root information filter (SRIF) to perform OD. The probe states are estimated in Cartesian coordinates, and the simulated time frame is 24 hours. The filter solution is mapped every 5 hours for post-processing analyses, as well as after each hopping event.

The set of estimated parameters and related uncertainties is shown in Table 1. A-priori covariances are assumed to be large, to model poor initial knowledge of the scenario. The mother spacecraft position is a consider parameter: its value is not estimated, but its uncertainty is considered in the estimation process. While in reality it would be possible to co-estimate the spacecraft state, we make this conservative assumption here to compensate for the lack of modeling of the spacecraft dynamics. A-priori state uncertainties are always spherical. The probes' state uncertainties are re-initialized every time a hop occurs, to account for its highly random nature.

3.4 Gravimetry Analysis

In this section, the most relevant analyses and considerations related to gravity estimation are summarized. First, the overall gravimetry performance for the baseline scenario is presented and discussed. Second, the sensitivity of the gravity estimates to some key design parameters is evaluated, in order to identify bottlenecks in the mission architecture. Such sensitivity analyses are performed using one single probe trajectory, spanning a 24-hour time frame, and using a ground-truth gravity field truncated at degree and order 20. In this case, we compute the full OD fit (i.e., not only a covariance analysis) in order to compare the residuals between the different cases analyzed. The resulting trajectory is shown in Figure 7.

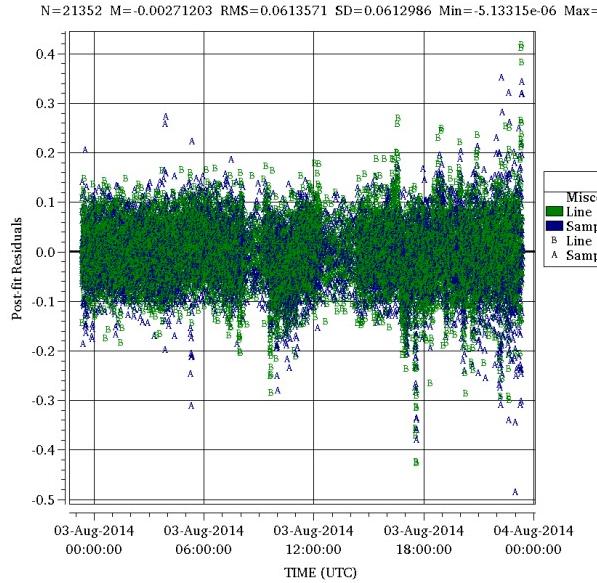
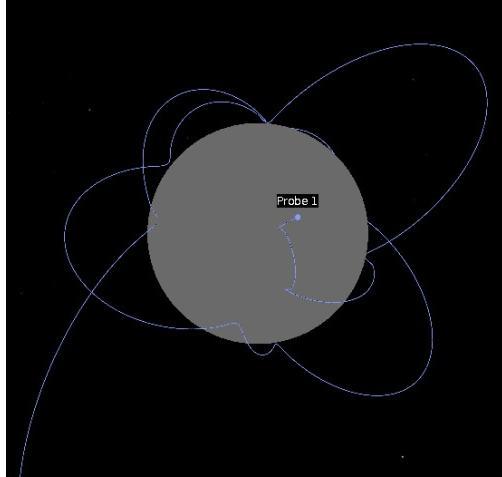


Figure 7: Left: Sample trajectory used for sensitivity studies. Right: postfit residuals of the baseline sensitivity case, at the zero-th filter iteration.

3.4.1 Baseline Covariance Analysis

A covariance analysis of the baseline design described in Section 3.3 is presented. The objective is to quantify the Signal-to-Noise Ratio (SNR) of the estimated spherical harmonic coefficients, i.e., the ratio between the a-priori and the a-posteriori covariances. The probe trajectories used for gravity estimation are shown in Figure 8 (left), plotted in a body-fixed frame. The resulting spherical harmonic observability is shown in Figure 8 (right), where it is compared with OSIRIS-REx gravity estimates [26] (both with the use of particle tracking and via standard spacecraft OD). It should be noted that the gravity field observed from OSIRIS-REx does not coincide with the one in our simulation, which is procedurally generated; however, the spherical harmonics are in the same order of magnitude for each degree and follow the same power law.

Figure 8 clearly shows that the Gravity Poppers scenario provides a gravity SNR at least two orders of magnitude higher than in the OSIRIS-REx case: not only is the SNR offset at degree 2 about two orders of magnitude higher, but it can also be deduced from visual inspection that the mean SNR decay rate (i.e. a log-linear trend) is approximately 4 times lower than for OSIRIS-REx. In the Gravity Poppers case, the SNR decays by about by 10^2 over 40 degrees, whereas the same decay is experienced by OSIRIS-REx data over just 10 degrees. Thus, we can extrapolate these results to infer that this particular simulation setup is capable of resolving harmonics (at $\text{SNR} > 3$) up to degree 80. Future work will verify this SNR decay trend using higher-degree gravity models and accounting for additional dynamics effects, such as the TRP.

In summary, this covariance study of the baseline mission design demonstrated that the Gravity Poppers concept has great potential for high-fidelity gravity estimation, which far exceeds state-of-the-art techniques. Nevertheless, this scenario is still based on a mixture of conservative assumptions on the one hand, and simplified dynamics models (e.g., TRP and

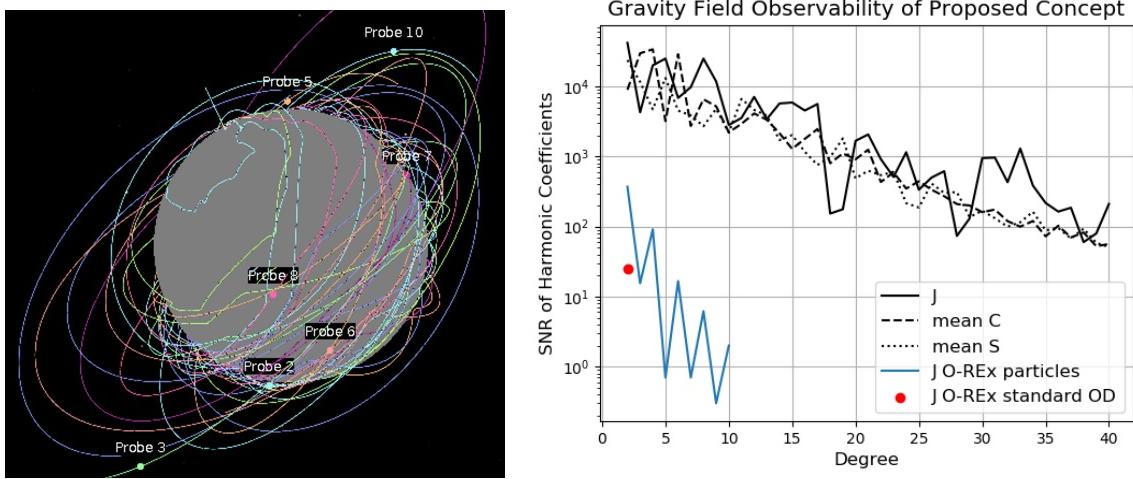


Figure 8: Baseline simulation results. **Left:** 3D visualization of the 10 probe trajectories. **Right:** SNR spectrum of the observed gravity field, compared to the OSIRIS-REx results.

albedo pressure) on the other hand. Thus, future work will more accurately characterize strengths and limitations of this approach, using more advanced simulations and computational capabilities.

3.4.2 Sensitivity to Imaging Frequency

The first parameter we showcase is the imaging frequency. When estimating small local perturbations on the dynamics, such as those given by higher-order gravity terms, the tracking frequency defines the temporal (and thus *spatial*) resolution of the measurements, which is essential to constrain the orbit determination solution and decrease uncertainties. Here, we compare two imaging frequency cases: a 10-second imaging interval and a 10-minute imaging interval. The former is in line with case study presented in the previous section, whereas the latter is similar to the interval used to track natural particles during the OSIRIS-REx mission. The trajectories and overall simulation setup are the same for both cases, i.e., only imaging frequency is varied. The resulting gravity estimates are shown in Figure 9: the higher imaging frequency provides an SNR improvement of almost 5 orders of magnitude! Note that the 10-minute imaging interval provides estimates in line with the OSIRIS-REx case (which also used approximately the same temporal resolution in tracking).

We claim that the quantification the immense value provided by higher tracking frequencies on gravity observability is one of the key lessons learned from this study. Such a high temporal resolution is only possible using onboard image processing to detect probes across frames, as link limitations would not allow downloading the immense data volume of such frequent images. Thus, onboard probe detection is necessary to enable high-precision gravimetry. As discussed in Section 4, the design of bright probes (relative to the small body and the dark background), as seen from the mother spacecraft, ensures that onboard detection can be performed robustly, efficiently, and with high centroiding performance.

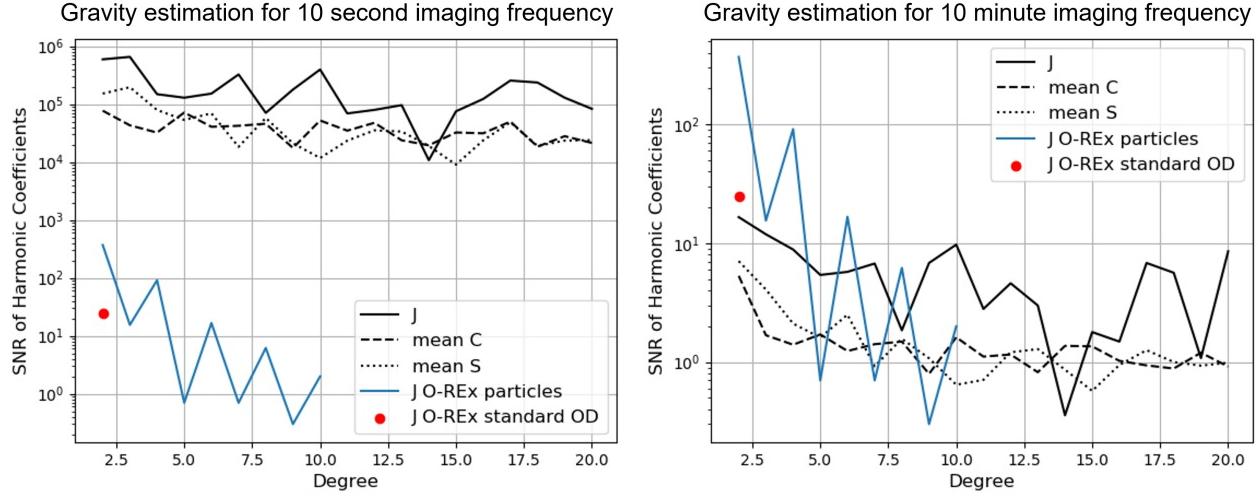


Figure 9: Spherical harmonics spectrum for two different imaging frequencies: 10 seconds (left) and 10 minutes (right).

3.4.3 Sensitivity to Probe Centroiding Precision

The second key parameter related to probe tracking is the uncertainty associated with probe detections. Usually, this task is executed by a detection algorithm that performs background subtraction (if needed) and eventually centroiding. For sub-pixel objects (such as the probes in this mission concept) centroiding involves determining the center of the probe's point spread function. Hence, the accuracy of detection increases with the probe's photometric SNR in the image. When the point spread function of a sub-pixel object is larger than 1 pixel, and the SNR is sufficient for detection (e.g., $\text{SNR} > 3$), centroiding error is sub-pixel [32]. For example, as shown in Section 5.3, the Gravity Poppers can be designed to be quite bright relative to the small body, hence the centroiding error is expected to be very low (about 0.01 pixels). However, here we relax the assumption of high-precision centroiding and assess how this impacts gravity estimates. Figure 10 compares the spherical harmonic SNR spectrum obtained from the baseline centroiding uncertainty of 0.05 pixel with the SNR obtained using a conservative 1-pixel uncertainty (a worst-case scenario). The gravity SNR is about one order of magnitude lower for the 1-pixel centroiding error. Although this can seem as a substantial decrease, it is small compared to the SNR magnitude, which is in the order of 10^6 , for low degrees. Moreover, the 1-pixel uncertainty does not undermine gravity observability, as the spherical harmonic terms remain largely observable, i.e., $\text{SNR} \gg 3$ for all gravity degrees considered here.

Lower-precision centroiding is usually related to observations with lower SNR, which can be caused by probes that are dimmer, further away, or in front of the lit small body. Future sensitivity studies will need to further characterize the impact of centroiding errors out to higher degree terms. However, this initial study suggests that better centroiding does help, but is perhaps not a primary concern. We speculate that lower tracking accuracies can be compensated for, to a large degree, by higher measurement *frequencies*.

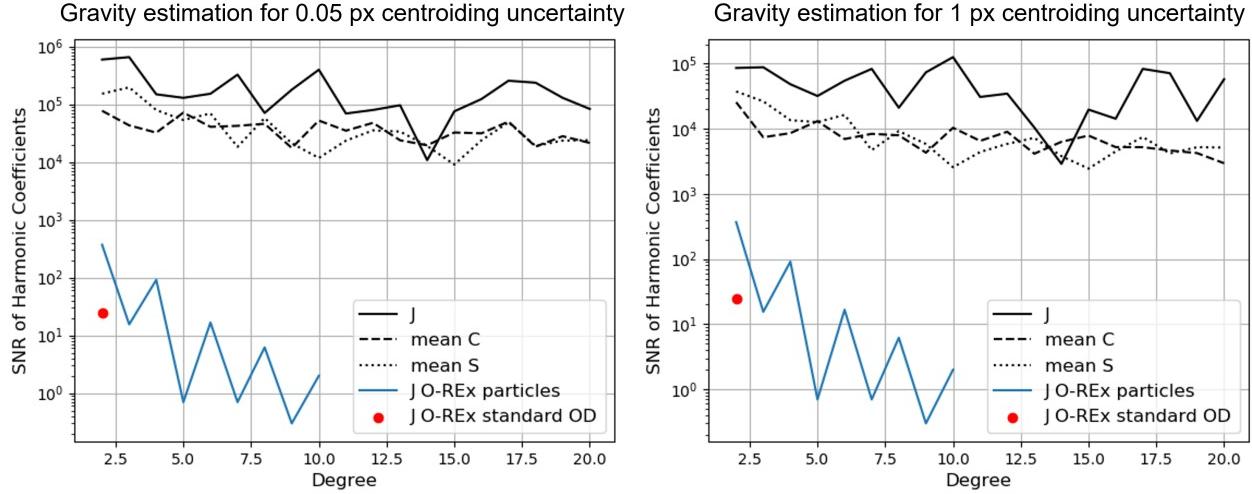


Figure 10: Spherical harmonics spectrum for two different measurement uncertainties: 0.05 pixel (left) and 1 pixel (right).

3.4.4 Ranging Measurements

One of the objectives for this Phase I study is to identify the best probe tracking strategy, i.e., the type of tracking measurements which maximize gravity observability. To contain mission complexity and overall budget, we claim that the most compelling observation types are optical (i.e., angle) and radiometric (i.e., range and range-rate) measurements. As shown by the above results, optical measurements have a great potential for this mission architecture, and can substantially outperform state-of-the-art gravimetry, in and of themselves. The key question we address here is whether the addition of radiometric measurements to optical tracking can contribute to further decrease the state (and gravity) uncertainties. To answer this question, it is important to quantify the tracking precision provided by optical-only measurements to understand its limitations, in the first place. In this regard, Figure 11 shows the uncertainty RMS evolution related to the position and velocity of one example probe. The position uncertainty drops below cm-level after 10 hours of observation, and the velocity decreases to hundredth of mm/s within the same time frame. To bring a substantial improvement to such covariance values, the radiometric measurements should achieve tracking precision which is at least in the same order of magnitude of these state uncertainties. We validated this conjecture with a simulated case study: we implemented a notional ranging system in our software pipeline, which performs ranging with a measurements uncertainty $\sigma_{msr} = 10\text{ cm/s}$, with a tracking interval of 60 seconds. Under such assumptions, no improvements in gravity SNR and probe state estimation could be observed, as expected.

Given the substantial increase in mission-design complexity associated with implementing a radiometric tracking system, and the fact that optical measurements, in combination with a high-precision dynamics model, provides excellent tracking performance, we conclude that optical-only tracking is the ideal choice for measurement strategy. Gravimetry performance may be further increased by optimizing key design parameters such as the probe hopping dynamics, or simply performing longer observation campaigns, while still relying on optical-only measurements for orbit determination.

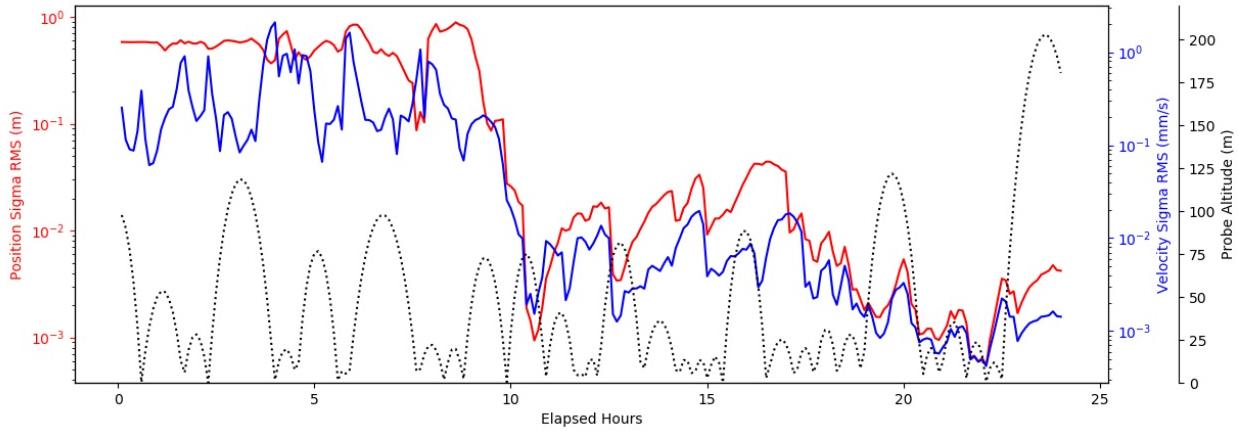


Figure 11: Evolution of the position and velocity uncertainty (RMS) for one hopping probe using optical tracking only, compared to their hopping altitude.

3.4.5 Boulder Weighing Case Study

The Gravity Poppers mission concept has the potential to provide invaluable insights about the mass distribution underneath and across the body surface. While spherical harmonics are a powerful model to capture macroscopic distribution patterns that apply to the whole body interior, they are not best suited to model a concentrated mass (a so called *mascon*) which affects the gravity field. Mascons are common features of small bodies: large boulders on the surface, craters, and subsurface voids are all examples of elements that can accurately be represented by a mascon. (For craters and voids, the mascon value would be negative.) Furthermore, the estimation of such parameters can provide additional insights for planetary science. For example, estimating the mass of a boulder on the surface of a rubble pile asteroid can be used to compare the boulder’s density with the asteroid’s bulk density, which in turn can be used to infer the internal structure and porosity.

We tested this “boulder weighing” concept in a simulated case study, again using asteroid Bennu, which makes sense given it is an apparent rubble pile with many surface boulders. We use the same ground-truth spherical harmonic gravity model as in the previous section, but this time, we add a mascon to the surface with a gravitational parameter of $\mu = 0.0087 \text{ m}^3/\text{s}^2$, which roughly corresponds to the expected mass of BenBen Saxum, the largest boulder on the surface of Bennu (approximately one thousandth of Bennu’s mass). We place this mascon at an arbitrary position corresponding to 45 degrees in latitude and longitude. The observability is evaluated on a 24-hour observation campaign, using the baseline covariance analysis framework previously presented. Results are shown in Figure 12 and compared to the observability of zonal harmonic coefficients. The temporal evolution of the mascon’s SNR is similar to the harmonic coefficients; interestingly it is very similar to the trend for J_{22} which represents an effective “wavelength” that is similar to the size of BenBen. Since the SNR at the end of the 24-hr observation is approximately 10^4 , we can say that the density of a boulder the size of BenBen may be estimated to better than 0.1%!

This case study demonstrates that the Gravity Poppers concept is capable of answering a meaningful science question related to the mass distribution of a rubble pile asteroid.

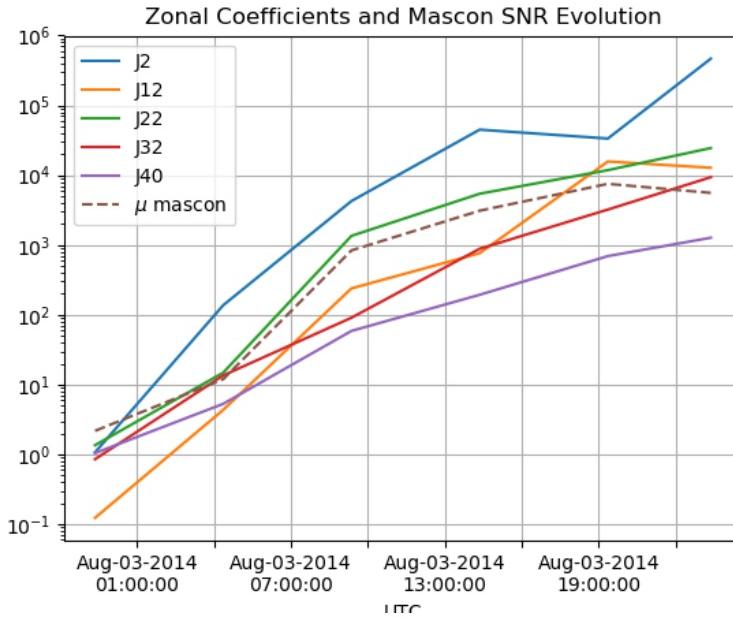


Figure 12: SNR evolution for the zonal coefficients and the mascon representing a surface boulder, using the baseline covariance study.

However, this simple, single-boulder weighing case study just scratched the surface; the information content in the high-order harmonic gravity fields revealed in this study can likely be inverted to estimate much higher resolution mass distributions than a single boulder. Future work will need to extend the gravity field representation of the body to a model that applies to highly irregular shapes, one which directly estimates mass distributions, or a hybrid of the two.

3.5 Numerical Considerations

One key finding of this Phase I study that was not expected, is the sheer computational challenges associated with the estimation of very deep gravity fields. The estimation problem, as conceived, is essentially a very large batch least squares filter, which requires the inversion (or pseudo-inversion) of a very large matrix, which grows as the number of measurements squared. Thus, we ran into computational bottlenecks for gravity fields beyond about degree 40 and for simulations lasting longer than 24 hours. Because of this, we were not able to fully characterize the gravity field observability out to the noise floor for our baseline scenario (out to degree 80), nor were we able to sufficiently address the questions of *how many probes* there should be, *how long the probes should be observed for*, and what is the point of diminishing returns for *increased imaging frequency*, all of which require handling much larger data sets. However, it is clear from figure 12 that the SNRs do not yet plateau at the end of a 24-hour simulation, suggesting that deeper gravity fields may be resolved with longer observation times. Future studies should consider using substantially increased computation resources (speed, parallelization, memory, and data storage) as well as estimation algorithms that make more efficient use of data in order to fully exploit the vast data sets possible through this mission concept.

4 Tracking System

The second key focus of our Phase I study was the tracking system, whose purpose is to generate observations of the probes' motions to be fed into the gravity estimation and orbit determination algorithm. As discovered through sensitivity studies (discussed in Sect. 3.4), gravity observability—especially for the weaker terms—is sensitive to the tracking accuracy and especially frequency. Thus, it is crucial to develop tracking strategies that are very precise and robust. This section summarizes the trade space of sensor architectures considered for probe tracking (Sect. 4.1), and then focuses on the optical tracking system (Sect. 4.2), establishing the basic principle of optical detection and demonstrating autonomous multi-probe tracking on simulated imagery.

4.1 Tracking Sensors Trade Study

The methods by which the trajectories of hopping probes can be measured was a key trade considered in Phase I, as it is the primary input to the gravity estimation filter. Various sensor architectures were considered, including both position and velocity sensors such as cameras, Lidar, and Doppler/ranging radios. Accordingly, three measurement modalities were considered: (1) angles-only from the mother spacecraft, (2) Doppler/ranging between the mother spacecraft and each probe, and (3) Doppler/ranging between all pairs of (local) probes. These modalities are illustrated in Fig. 13. Different combinations of these modalities represent trade-offs of performance and complexity.

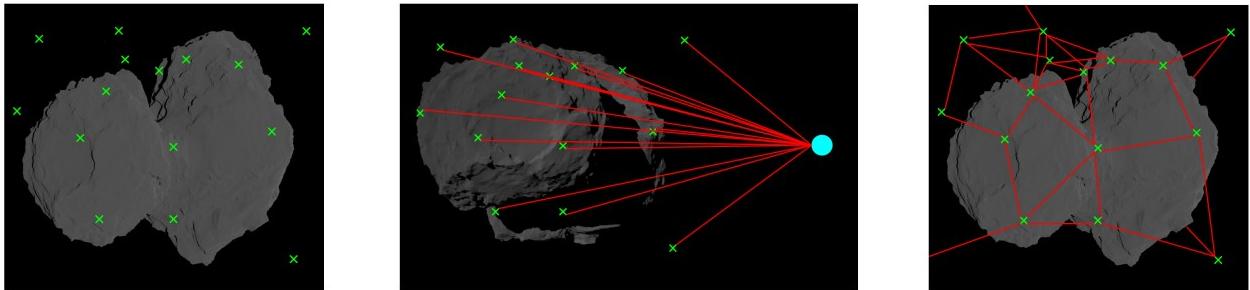


Figure 13: Three tracking modalities considered in this Phase I study: angles-only from the mother spacecraft (left), Doppler/ranging between the mother spacecraft and each probe (center), and Doppler/ranging between all pairs of (local) probes (right).

An angles-only modality (Fig. 13-Left) can be achieved with either an onboard camera or laser scanner (Lidar). A camera requires that the probes be bright enough in the sensitivity spectrum and can achieve centroiding accuracies down to 0.01 pixels, or about **1 μ rad** for a 10 deg FoV, 4 Mpx camera. A Lidar requires the probes to be highly retroreflective, and generally has poorer angular resolution (**100 μ rad** for OSIRIS-REx's OLA instrument [33]). Range/range-rate measurements to the probes (Fig. 13-Center) can be achieved with specialized radio communication protocols or Lidars. The ranging precision of time-of-flight radios can vary considerably depending on the bandwidth, clock stability, and other factors, but **10 cm** is a common value achieved by UWB ranging sensors. Doppler measurement precision can also vary considerably, but commonly achieve **0.05 - 1 mm/s**. OSIRIS-REx OLA instrument achieves about **30 cm** range precision. Finally, ranging *between* probes would require omnidirectional transmission and reception of radio signals, as probes are not

assumed to be attitude controlled. However, inter-probe ranging would require a more sophisticated communication protocol and their maximum range is likely limited due to their low power budget. Thus, it is likely not practical for this mission concept.

Both angles and range/Doppler measurements were considered in our gravity estimation sensitivity studies in Sect. 3.4. We concluded that *optical measurements* (i.e. probes imaged by the mother spacecraft) are sufficient to characterize the small body’s gravity field with high precision and that range/Doppler measurements did not noticeably improve performance. Even though optical measurements do not themselves constrain the depth dimension, when processed in a batch estimation filter, the post-fit residuals on the probe position estimates were on the order of 1 mm - 1 cm, suggesting that ranging sensors would likely need to have a similar (or better) precision to appreciably reduce uncertainties. Similarly, post-fit velocity residuals ranged from 0.1 - 0.001 mm/s, suggesting that Doppler precision would need to be on that order to further reduce uncertainties. Thus, for the rest of this study, we focused on Optical measurements, as they provide very high tracking accuracies with little complexity.

4.2 Optical Tracking

Optical measurements were determined to be the most promising candidate method for probe tracking. The basic principle of optical tracking is illustrated in Fig. 14, with key parameters annotated. In this section we discuss the key aspects of an optical tracking system: the camera selection, imaging geometry, probe detection, and autonomous multi-object identification and tracking.

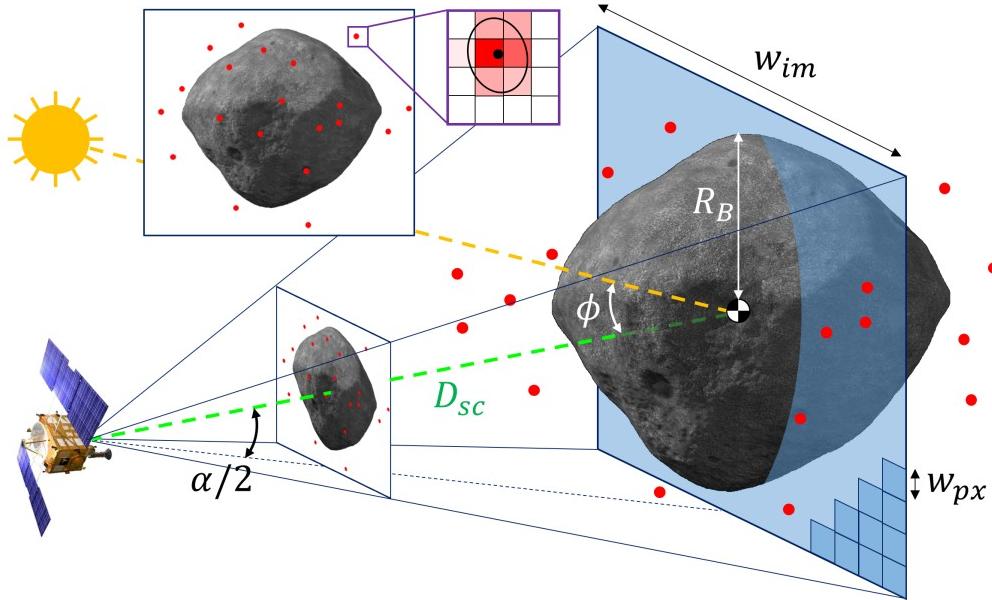


Figure 14: A schematic illustrating the basic concept and geometry of optical tracking, with key dimensions annotated.

4.2.1 Camera Selection

A huge variety of cameras have been developed for in-space imaging applications, ranging all the way from cell-phone grade context cameras to multi-billion-dollar telescopes such as

Hubble. In general, for increased photometric sensitivities and astrometric accuracies, size matters. Larger apertures, larger sensors, and an increased number of pixels are all beneficial. As a baseline, we select a state-of-the-art astrometric-quality camera demonstrated on the “Arcsecond Space Telescope Enabling Research in Astrophysics (ASTERIA)” CubeSat [31], a 2U camera that is both high performance and small enough to be hosted on a number of spacecraft. Table 2 captures the ASTERIA’s camera properties of interest for this analysis.

Parameter	Units	Value
Aperture diameter	mm	60.7
Focal length	mm	85
Detector dimensions	pixels	2592×2192
Pixel size	μm	6.5×6.5
Detector Field of View	deg	11.2×9.6
Passband	nm	500-900
Quantum Efficiency in Passband	–	0.42
Lens Throughput	–	0.8
Gain	e^-/ADU	6.44
Bit Depth	–	11
Assumed Parameters	Units	Value
Read Noise	e^-/pixel	15
Dark Current	$e^-/\text{pixel}/\text{s}$	0
Integration Time	s	1.5
Subtended Solid Angle	sr	$1.8 \cdot 10^{-10}$

Table 2: ASTERIA’s camera specifications

4.2.2 Spacecraft Distance

A key parameter that needs to be defined for the optical tracking analysis is the *distance* of the spacecraft from the body (D_{SC}). The distance, in combination with the field-of-view, α (set at 10 deg according to table 2), determines the physical span of the small body in the image (w_{im}). For this mission concept, there is an inherent trade-off in setting an appropriate span: too close and only the probes in the immediate vicinity of the body are visible, but too far away and the resolution of the probes and hence their SNR decreases, negatively affecting astrometry performance and hence gravity estimation. Thus, we chose a nominal image span such that the small body occupies about 70% of the frame width, which sets the spacecraft at a nominal distance of 4 km for our Bennu case study. Given the small gravitational attraction of a small body such as Bennu, this distance is a good candidate for an inertial hovering orbit, which would be ideal to perform particle tracking. Hence, this simplified analysis demonstrates the feasibility of realizing a mission geometry suitable for probe tracking.

4.2.3 Probe Detection

Now that a notional camera has been selected and the spacecraft imaging distance has been set, we can use a photometric analysis to determine how *bright* the probes need to be in order to be detected. As discussed in Sect. 3, the ideal case that we will consider here is that

the probes are bright enough to be detectable while transiting in front of the small body. Obviously, if it is proven that the probes are visible in front of the small body, the visibility requirement also fulfills detection over the dark sky background.

The first consideration for assessing probe visibility is to determine if the imaging system is diffraction-limited. The standard deviation of the point-spread-function (PSF) of a particular camera can be computed as,

$$\sigma_{PSF} = 0.45 \cdot \lambda_{max} \frac{f}{D_A p} = 0.09 \text{ pixel}, \quad (2)$$

where $\lambda_{max} = 900 \text{ nm}$ is the maximum wavelength within the passband (to account for the highest PSF standard deviation), f is the focal length, D_A is the camera's aperture diameter and p is again the pixel pitch. As shown, this value is much smaller than one pixel, indicating that images from this camera are limited by the sensor's resolution, rather than by diffraction. Since the physical span of the probe in the camera field of view is also generally smaller than one pixel, the photon flux from the probe is entirely contained within a single pixel, for a focused image. However, it can be shown that a one-pixel detection yields to very poor astrometry performance. One possible solution to this is to intentionally defocus the optical system, so that the PSF distributes over multiple neighboring pixels. Using defocused optics to improve performance is an established technique in astrometry [34]. The downside of defocusing is that, since photons distribute over multiple pixels, the intensity of the central pixel decreases, and so does the SNR. If it is assumed, for example, that defocusing yields $\sigma_{PSF} = 1 \text{ pixel}$, then only 68% (1- σ span) of the photons lie within the central pixel. This can be accounted for when sizing the probe by using a degradation factor of $\eta = 1/0.68$.

Finally, an expression accounting for all noise sources and the small-body background can be derived to compute the required probe brightness for a certain SNR¹:

$$SNR = \frac{\eta \Phi_P Q \tau}{\sqrt{\Phi_P Q \tau + N_{DC} \tau + N_R^2} + \Phi_B^* Q \tau}, \quad (3)$$

where Φ_P is the photon flux of the probe, Φ_B is the maximum expected photon flux of the small body given its reflectance properties, τ is the exposure time, Q is the camera efficiency, and N_{DC} and N_R are the dark-current and read noise, respectively. For a desired SNR for detection, this quadratic equation can be solved for Φ_P to determine the brightness required for detection. Section 5 will investigate two methods the probes may use for achieving this brightness requirement: through passive sunlight reflection and LED strobining.

Background Subtraction: Now assuming that probes achieve a sufficient SNR to be detected, the small body background may be “subtracted.” Many background subtraction techniques exist that should be investigated in future studies. Here, we use a simple technique (also used by the OSIRIS-REx team): a median filter, with a 15×15 window, which works best when operating on the dark-sky background, but can also help to reduce the noise of the small body image. An example of a background-subtracted image is shown in Figure 15.

¹It should be noted that the given expression assumes a CCD sensor, commonly used for scientific applications, whereas the ASTERIA camera is provided with a CMOS sensor. However, given the notional-design nature of this analysis, this does not undermine the overall conclusions.

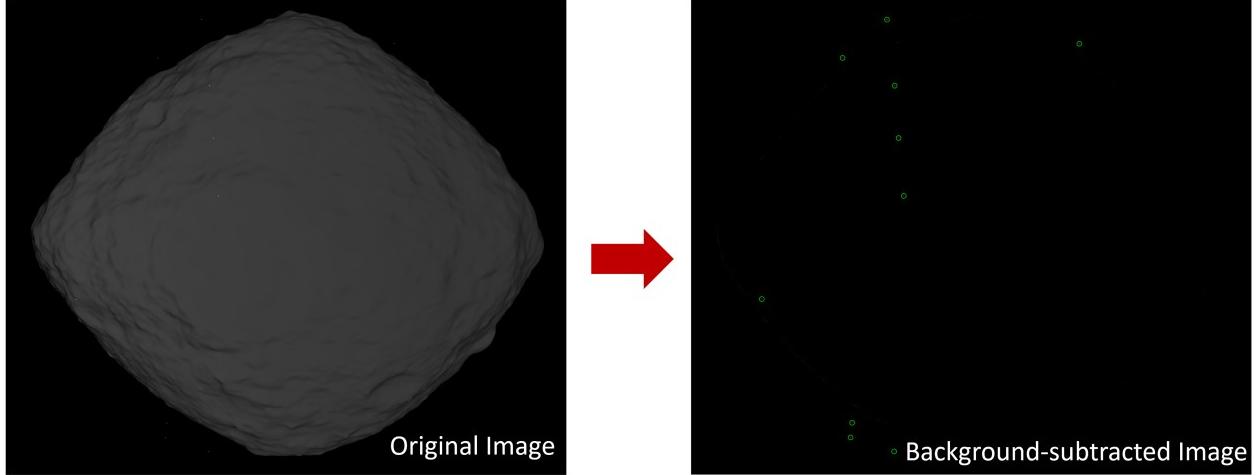


Figure 15: Example of background-subtracted synthetic image. **Left:** Original image. **Right:** background-subtracted image. The green dots are the centroids of the probe detections.

Centroiding: Once the background-subtracted image is available, it is possible to detect the probes and compute their centroids. The first step is to use thresholding to create a binary mask of all the pixels whose intensity exceeds the threshold (e.g. 50% of the maximum intensity in the image). The second step is to use a connected-component algorithm (e.g. MATLAB’s `bwconncomp` function) to define the pixel clusters for each probe. The key advantage of this clustering algorithm is that it does not require the number of clusters as an input (which is the case for *k-means*, for example) and the number of detected clusters is an output of the routine itself. These clusters do not capture the dimmer pixels related to probes PSF, but their centroid can be used as an initial guess for the probe centroid. Therefore, the third step is to center a tracking window around each of the cluster centroids. The pixel size of the window is defined as $W = 4\sigma_{PSF,px}$. An iterative process is used to compute the moment-based centroid within the tracking window and re-center the window around the new centroid, until the centroid solution converges. The centroids resulting from this process are considered detections for the subsequent multi-object tracking process. The astrometric uncertainty is also computed and is below $\sigma_C = 0.02$ pixel for the example in Fig. 15.

4.2.4 Multi Object Tracking

The last and crucial step in the pipeline is the multi-object tracking process, which takes time-stamped unidentified detections as input, and identifies a set of tracks as an output. This is done associating detections with tracks, using Multi-Object Tracking. A Track-Oriented Multiple-Hypothesis Tracking (TOMHT) approach is used here, leveraging the MATLAB’s `trackerTOMHT` function. All detections are considered here as confirmed measurements: given the high SNR of the probes with respect to the background, the probability of false alarm is considered negligible. This increases the robustness of the whole tracking system.

The dynamics propagation is based on a constant-velocity linear Kalman Filter. The use of such a simple dynamics model is justified by the relatively high imaging frequency: if the time interval between frames is small enough, higher-order dynamics effects can be neglected. The downside of this approach is that it requires continuous tracking of the

probes to guarantee good performance, which is not the case when, for example, the probe trajectories are occluded by the small body, the probes move out of frame, or when off-nominal conditions interrupt imaging sequences. To overcome these challenges, future work should implement a non-linear Kalman Filter, e.g. an Extended Kalman Filter, with a two-body Keplerian dynamics perturbed by a solar-radiation-pressure acceleration, as shown by Liounis et al. [35]. With these filters, long interruptions between probe detections could be tolerated and correctly process, allowing, for example, a probe to be tracked even if it is temporarily eclipsed by the body.

The TOMHT algorithm is based on several tunable parameters, which highly affect overall performance. For this study, we tuned the parameters based on knowledge of the scenario and some iterative refinement. The set of the key parameters tuned for this test are shown in Table 3. Future work should develop a more rigorous approach for automatic parameter tuning, and stress-test such parameters on several scenarios of interest.

Parameter	Value
Maximum number of tracks	10
Maximum number of track branches	10
Assignment threshold, $[C_1, C_2, C_3, C_4]$	$200 \cdot [4, 5, 12, 14]$
Probability of detection	0.99
False alarm rate	10^{-9}
Rate of new target per unit volume	10^{-14}
Maximum history steps considered	10
Probe state <i>a-priori</i> covariance	1 pixel (pos.), 50 pixel (vel.)

Table 3: Tuned Parameters for the TOMHT probe tracking algorithm.

A simulation was run to test the TOMHT performance. The algorithm processed 80 frames, with a time interval of 100 seconds per frame. This imaging frequency is likely slower than the imaging frequency that would be used in practice, and thus represents a challenge case for tracking, as the probes are displaced several pixels (5-20 pixels) from one frame to the next.

An example of TOMHT performance is shown in Figure 16, plotting the accepted tracks whose life is longer than 4 frames. As the figure shows, all the probes are detected and continuously tracked, in front of both the body and the black background, and their trajectories match the ground truth. The tracks also show robustness to the discontinuities occurring when the probes hop from the surface; this is an important feature in tracking, given the periodic bouncing behavior of the probes. The constant-velocity linear Kalman Filter is mostly able to provide good frame-to-frame estimates. However, there are a few pathological cases where the tracking algorithm fails: when the trajectory turn rate becomes too large (e.g. in the top-left orange track, after the surface hop) and at some intersections between tracks (e.g. at the intersection between blue and red tracks, on the right). These phenomena are likely symptoms of poor dynamics estimation for such challenging scenarios. As previously mentioned, future work will implement a more complex dynamics model, as well as a more advanced tracking pipeline, e.g. employing the “Mahalanobis distance” as a metric for measurement association. These upgrades are likely to resolve most of the tracking issues related to modeling the probe dynamics.

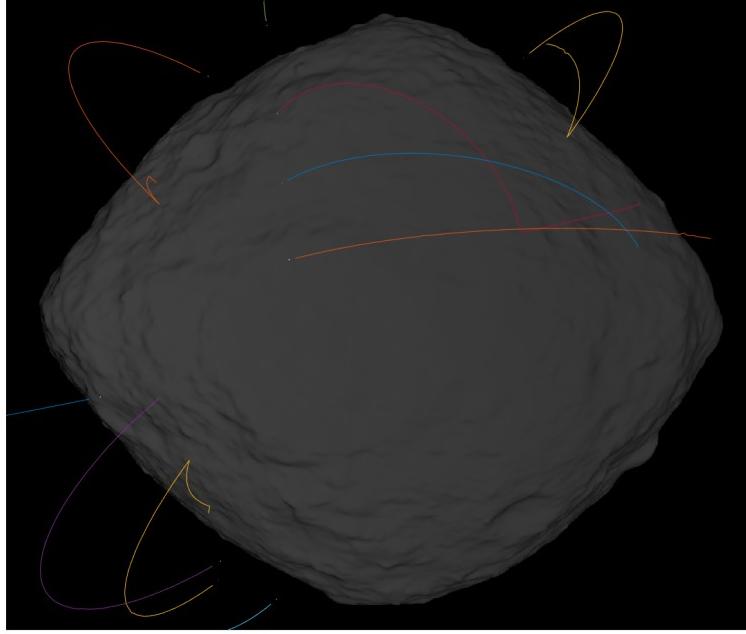


Figure 16: Demonstration of TOMHT multi-object tracking algorithm successfully identifying and tracking 10 probes simultaneously.

5 Probe Design

So what exactly is a “gravity popper?” What makes them effective “gravity beacons” and how does that inform their design and operation? This was the third key focus of the Phase I study. Our approach was to make no prior assumptions about the nature of their design or operation, but rather to allow their form and function to derive from the requirements and performance sensitivities established through our studies of gravity observability and tracking systems. This design philosophy is especially important for such a complex and multidimensional problem, where initial design intuitions are easily (and frequently) broken.

This section describes the probe design process of this Phase I study, from a blank piece of paper to a notional design concept. We start by summarizing the key requirements and performance metrics derived from our studies of gravity observability and tracking (Sect. 5.1). We then focus on the two most critical requirements and present the design trade space for (1) their optical visibility to a camera onboard the mother spacecraft (Sect. 5.3), and (2) their hopping dynamics and operation (Sect. 5.2). We then discuss two key subsystems of the probes: power (Sect. 5.4) and thermal (Sect. 5.5), which are common challenges for miniaturized space systems. We present two notional probe design architectures in Sect. 5.6 that differ in their method for achieving the brightness requirement for optical detection: one which uses passive sunlight reflection and one which uses LED’s to emit light. We further detail the system engineering of the LED-strobing probe, demonstrating first-order engineering feasibility of key subsystems and its notional integration into a deployment canister and tracking system.

5.1 Probe Requirements

Fundamentally, the purpose of the gravity poppers is to act as dummy masses whose gravitational forces from the small body can be measured through their ballistic motion. From our analyses in Sects. 3 and 4, we can translate this objective into two fundamental requirements:

- **Hopping:** The probes shall be capable of hopping at suborbital speeds from any location on the body and on any likely terrain properties.
- **Brightness:** The probes shall be bright enough to be detected from a distant camera, even when in front of the illuminated body.

These requirements, and the design trade space considered to meet them, will be discussed in detail in Sects. 5.2 and 5.3. However, there are additional performance metrics that should also be considered in the probe design:

- **Size (mass/volume):** The smaller the probes can be made, the more probes can be accommodated in a given mass/volume allocation, or conversely, the smaller and cheaper the overall instrument can be, thus increasing the possibility of mission infusion. Furthermore, higher *density* probes are also favorable for their higher ballistic coefficient, thus being more sensitive to gravitational forces relative to photon disturbances.
- **Simplicity / Robustness:** Complex probes and subsystems may enhance capabilities, but often at the cost of robustness, as there are more opportunities for failure. Wherever possible, the simpler and more robust solutions should be adopted.
- **Shape symmetry:** The probes benefit from being symmetric so that photon pressures are more consistent and predictable in any arbitrary rotation state.
- **Longevity:** A probe that hops for longer produces more data for estimating higher resolution gravity fields. In practice, a nominal time frame will be defined for probe tracking, but the longer a probe can remain hopping, the more flexibility and margin this provides for the mission.
- **Impact Elasticity:** Unlike prior hopping rover concepts in which subsequent bouncing after hopping is considered a nuisance to be avoided, for this concept, highly elastic bouncing is favorable, as it leads to more ballistic observations for a single hop.
- **Hop Controllability:** All else equal, the ability to both tune the nominal speed of a hop and reduce its variability produces a more controllable and predictable swarm dispersion for improving gravity observability.
- **Communication:** The ability to both send commands to, and receive signals from the probes opens up possibilities for enhanced functionalities (such as probe identification) which may offset the added complexity of a small radio.

This list of key design considerations address the performance, risks, and cost associated with this mission concept, but they are often conflicting. For example, enhancing capabilities

such as hop controllability and radio communication typically come at the cost of size and complexity. Thus, these requirements and performance metrics are not likely to point towards a single solution, but rather guide the trade space of design options to consider. This Phase I study did not treat the probe design problem as a rigorous optimization study, but rather addressed the key functional subsystems in order to derive a *notional* probe architecture that meets the requirements of the mission.

5.2 Hopping System

In order to achieve the low altitude trajectories and large spatial distributions required to measure a high-resolution gravity field, collisions with the body are inevitable, and thus, a mechanism for *hopping* back into orbit is required. Luckily, hopping, even at orbital speeds, is easily achieved in the microgravity environment of small bodies, where gravity is generally between 10,000 and 100,000 times lower than on Earth. Indeed, many hopping mobility systems have been conceived for the exploration of small bodies [36, 37, 38, 39, 40, 41, 42, 43, 44, 45, 46, 47, 48, 49], two of which have actually been deployed to the surface of asteroid Ryugu from Hayabusa-2: MINERVA [46] and MASCOT [47]. Even the Philae lander deployed from Rosetta, which had three mechanisms for anchoring to the surface of Comet 67P, inadvertently hopped over 1 km after touchdown [50, 51]. So hopping is a very natural tendency for objects on small bodies. In this section we will derive the hopping speed requirement as a function of body size, present a trade space of simple hopping mechanisms, and discuss the operational logic for choosing *when* to hop.

5.2.1 Hopping Speed

Due to their simple design and operation, Gravity Poppers are assumed to have no knowledge of their orientation on the surface nor any control over the *direction* of hops (and subsequent bounces). However, they do have control over the input energy to the hopping actuator, thus giving them some degree of *speed* control. So how fast should the probes hop to maximize gravity observability? Too slow and they may not sufficiently redistribute around the body, getting trapped in local “geopotential basins” (e.g. a crater). But too fast and they may reach very high altitudes not useful for gravity observability, or even escape the body entirely. As a starting point, let’s consider the nominal escape velocity from the surface of a (non-rotating) spherical body:

$$v_{esc} = R_B \sqrt{8/3\pi G\rho_B}, \quad (4)$$

where R_B and ρ_B are the body’s radius and density, respectively. For a small asteroid like Bennu ($R_B = 250$ m, $\rho_B = 1200$ kg/m³), $v_{esc} = 20$ cm/s, and for a large dense asteroid like Eros ($R_B = 8.4$ km, $\rho_B = 2700$ kg/m³), $v_{esc} = 10$ m/s, which are energetically equivalent to free-fall heights of 2 mm and 5 m on Earth, respectively. However, asteroids are generally not spherical, and can often be rotating at significant rates. Thus, the true escape velocity depends on the specific location on the surface and velocity direction. To first order, a more realistic escape velocity can be defined as,

$$v_{esc}(\vec{x}) = \sqrt{\frac{2\mu}{||\vec{x}||}} - k_v ||\vec{\omega}_B \times \vec{x}||, \quad (5)$$

where μ is the gravitational parameter of the body, $\vec{\omega}_B$ is the body rotation vector, \vec{x} is the position on the surface (from the center of mass), and k_v is a factor ranging from -1, for a

purely westward hop, to 1 for a purely eastward hop. Thus, the escape velocity can vary significantly across the surface of a rotating irregular body. For Bennu, an eastward hop from the equator will reach escape at only 7 cm/s, requiring eight times lower hop energy than an escape trajectory from the poles. Accordingly, in order to guarantee that a probe will never escape the body, its hop speed could be set to not exceed the minimum escape velocity at any point on the surface. However, this speed constraint may be overly restrictive, limiting the altitude and duration of hopping trajectories and preventing effective redistribution around the surface. Indeed, the likelihood of a probe hopping in a purely eastward trajectory from the highest equatorial peak may be very low. Thus, in the paradigm of random hopping dynamics, we need a more *probabilistic* assessment of swarm spatial distribution and escape likelihood in order to select an appropriate hopping speed.

Probe Distribution: In order to select an appropriate nominal speed at which the probes should be designed to hop, a Monte Carlo dispersion analysis is required. A rigorous optimization study of probe hop speed to maximize gravity observability is a complex problem and left for future work. Here, we explore the trade-off between probe dispersion and likelihood of escape through a Monte Carlo case study for a hopping swarm on Comet 67P—a body with both a highly irregular shape and fast rotation, where $v_{esc,min} = 40$ cm/s and $v_{esc,max} = 118$ cm/s. Probes were initialized at random points on the surface and commanded to hop at a certain speed and in a random direction (within 60 degrees of the local normal vector). The ballistic dynamics were propagated with a reduced-order (5000-facet) constant-density polyhedral gravity model and the velocity was reinitialized in a pseudo-random bouncing direction upon collision with the surface, with a mean energy dissipation of 50%. Once the probe has settled on the surface, this process was repeated.

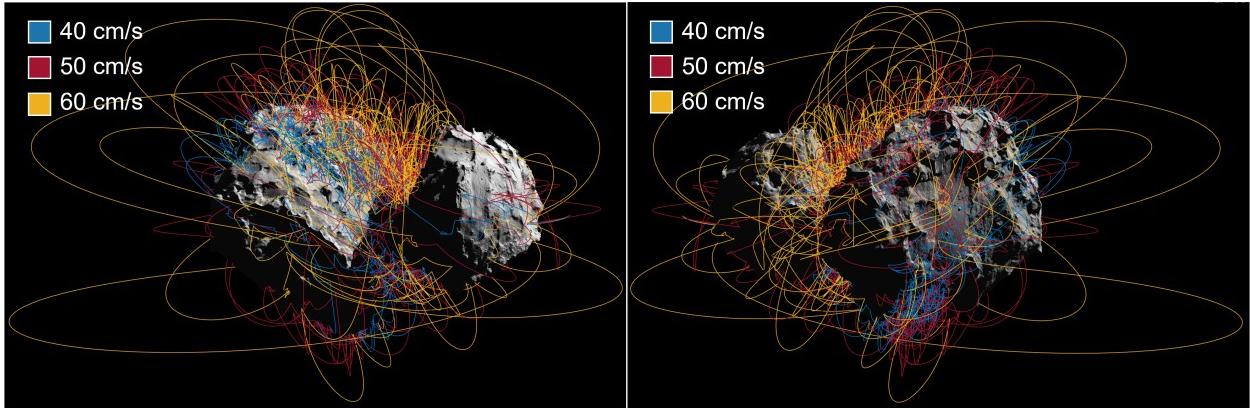


Figure 17: Monte Carlo simulation of three probes randomly hopping on Comet 67P for one month. Probes have preset hopping speeds of 40 cm/s, 50 cm/s, and 60 cm/s.

Figure 17 shows the trajectories of three probes on the surface of Comet 67P, simulated for one month. Each probe has a different preset hopping speed, one at 40 cm/s, one at 50 cm/s, and one at 60 cm/s. As expected, the trajectories generally exhibit a greater altitude and degree of redistribution at higher speeds. At speeds below about 30-40 cm/s, we found that probes tended to become “trapped” in the low geopotential neck region of the comet. The exact implications of such asymmetric spatial coverage (apparent in trajectories in Fig. 17)

on gravity field observability remains an open question that will require a more sophisticated analysis in future work. For now, we can say qualitatively that it is more desirable for the probes to be able to reach any point on the body from even the lowest geopotential basin, thus motivating more energetic hops. Thankfully, this simulation also shows that even hops at speeds above 40 cm/s—the theoretical maximum speed to *guarantee* no escape—the probes did not escape, suggesting that higher speeds may be practical for irregular bodies.

In order to quantify the probabilistic rate of escape as a function of hop speed, 10,000 probe simulations were run (1,000 at each of 10 different speeds) whereby the hopping dynamics were propagated indefinitely until the probe escaped the body. Figure 18 summarizes the results.

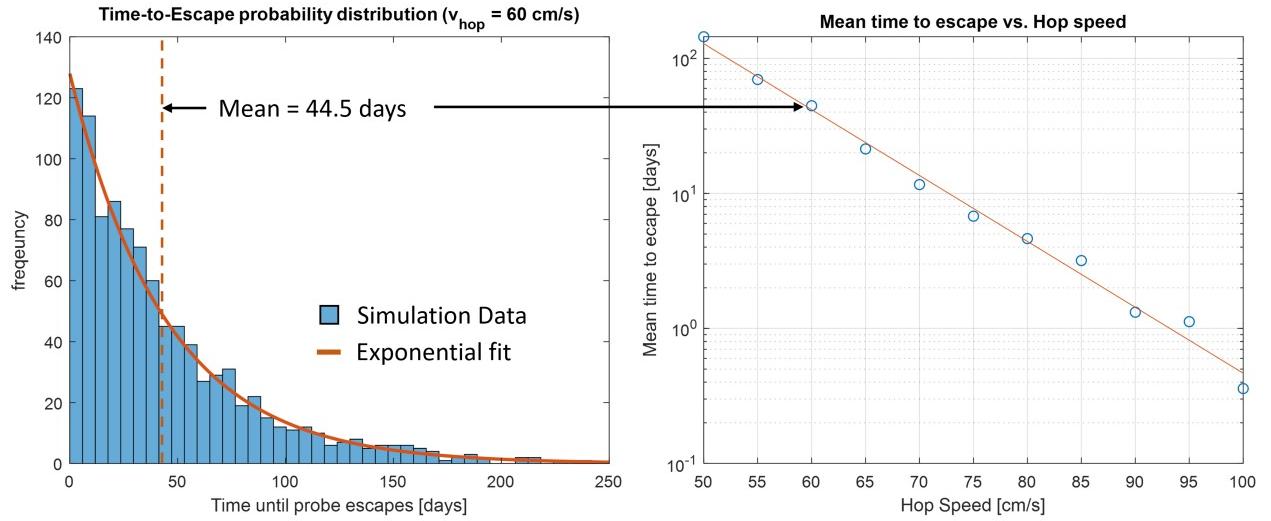


Figure 18: Statistical analysis of probe rate-of-escape when hopping on comet 67P with different speeds. **Left:** histogram of time-to-escape for Monte Carlo hopping simulations at 60 cm/s hop speed. An exponential fit is shown with a mean escape time of 44 days. **Right:** Mean time to escape for each batch of simulations run at different hopping speeds. The theoretical speed to guarantee no escape is $v_{\text{esc},\min} = 40 \text{ cm/s}$ and the theoretical speed at which escape is inevitable is $v_{\text{esc},\max} = 118 \text{ cm/s}$.

The results suggest that the time until a probe escapes the body is exponentially distributed, which makes sense given the Poisson random nature of hopping events. Also as expected, the mean time to reach escape increases for slower hops, theoretically approaching infinity as $v_{\text{hop}} \rightarrow v_{\text{esc},\min}$ and approaching zero (i.e. an immediate escape) as $v_{\text{hop}} \rightarrow v_{\text{esc},\max}$.

Now we have a tool that can be used to select the maximum hop speed of the probes for (1) an acceptable rate of escape and (2) a desired dynamic redistribution around the surface. Given that smaller hopping trajectories are provided “for free” during the subsequent bouncing of larger trajectories, a reasonable heuristic is to set the nominal hop speed as high as possible for an acceptable rate of escape. For example, for a two-week observation campaign in which it is desired to have at least half of the probes survive for the entire campaign, we can use the statistics of the exponential distributions characterized in Fig. 18 to select a nominal hop speed of 67 cm/s (corresponding to a *median* lifespan of 14 days). However, if the campaign is two months long, then the speed should not be set higher than 54 cm/s.

This analysis assumed that the target hop speed is able to be achieved exactly. Of course

control errors in the execution of hops and the energy absorbed by the uncertain surface properties may induce significant variability to this nominal target, as will be discussed in the next section. However, this type of Monte Carlo analysis can be easily extended to account for these uncertainties by treating the hop speed as a random variable rather than a fixed value. Future work should incorporate a model of the hop speed variability into this statistical analysis, which could be derived from experiments of the hopping system in a reduced-gravity testbed.

5.2.2 Hopping Mechanism

Now that we have a sense for the desired hopping *speed* as a function of body size, shape, and spin, how might the probes actually be designed to achieve these ΔV s when resting on the surface? Fortunately, many hopping robots have been developed both for terrestrial and planetary environments that we can draw inspiration from. Figure 19 highlights a handful of prior concepts that illustrate a spectrum of possible ways to achieve hopping, which can be broadly categorized into three main types: (1) internal momentum devices, (2) external appendages, and (3) propellant-based thrusting.

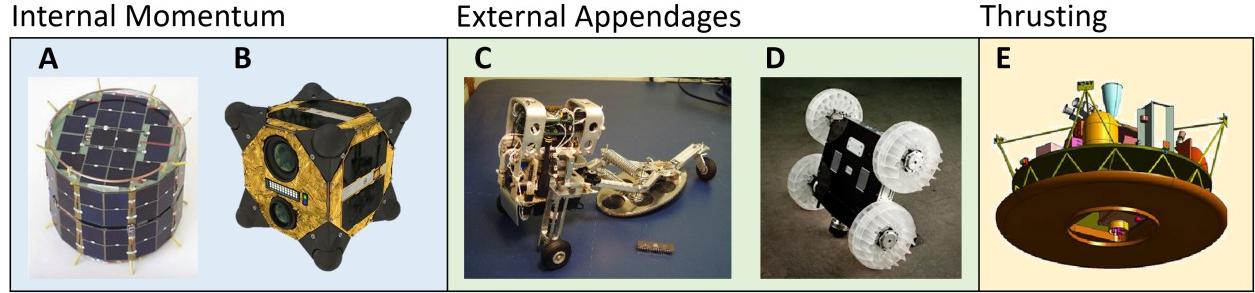


Figure 19: Examples of some hopping robots that leverage different principles of actuation. **A.** *MINERVA* rover uses a spinning internal eccentric flywheel to induce torques on the chassis and ground reaction forces [46], **B.** *Hedgehog* rover uses three (balanced) internal flywheels and mechanical brakes to hop via aggressive rotation [36], **C.** *Asteroid Hopper* uses an external spring-loaded “foot” to push on the surface for hopping [44], **D.** *Sand Flea* robot uses a very powerful compressed gas piston to hop at up to 10 m/s [52], **E.** *Comet Hopper* (CHopper) uses thrusters to hop [40].

Different types of hopping mechanisms have different advantages and disadvantages. Internal momentum actuators are attractive options for minimalistic probes due to their simplicity and omnidirectionality (i.e. able to operate in any orientation [53]), but are relatively inefficient and can be difficult to predict due to the rotational-to-translational momentum transfer through ground reaction forces. External appendages (e.g. spring-loaded foot pads) can be much more efficient and achieve higher speeds (up to 10 m/s in the case of Sand Flea [52]), but generally require the robot to be correctly oriented on the surface, thus requiring either passive or active self-righting capabilities. Thrusting is attractive for its consistency since it does not require “pushing” on the surface and is therefore insensitive to the uncertain surface strength. However, thrusters have inherent propellant limitations, they are directional (requiring pointing/self-righting), and they may be too complex for simple probes.

For our application of random hopping in microgravity, simplicity and robustness are the most important qualities for a minimalistic probe. Thus, internal momentum devices

are likely the best option since they have no external moving parts (and thus less likely to foul from dust or get caught on rocky terrain) and can hop in any orientation with a single actuator, as demonstrated in reduced gravity flights with the Hedgehog rover [53]. Flywheel-brake hopping systems are highly robust on irregular and soft terrain, and although they are fairly inefficient, hopping is so infrequent and requires such little mechanical energy that even an inefficient actuator is a negligible contribution to the overall power budget (as discussed in Sect. 5.4). Their sensitivity to the strength of the surface material can be mitigated to a large degree by incorporating “soft” elastic appendages that increase the stroke length of the impulse and thus, reduce the contact pressure. Once at rest on the surface, the flywheel would be slowly accelerated (so as to not induce rotation) and then abruptly stopped with a high-friction band brake (similar to Hedgehog [36]). This angular momentum transfer causes the chassis to abruptly rotate and press its elastic legs against the surface, inducing a hop. The speed of the resulting hop is directly proportional to the speed of the flywheel at the moment the brake is applied, thus allowing for simple speed control.

The flywheel-brake hopping mechanism has practical speed limits of about 1-3 m/s. Thus, if larger bodies like asteroid Eros or Mars’ moons Phobos and Deimos are considered for this mission concept, the probes will likely require a more energetic linear impulse mechanisms for hopping. And although our preliminary ballistic dispersion analyses suggest that the exact hop speed need not be precisely controlled, if this does become desirable then a simple cold gas thruster system may be considered, which could deliver a more consistent impulse since it is independent of the terrain. However, these options are likely to increase the mass and complexity of the probes.

One characteristic behavior of flywheel-brake hopping systems is their residual angular momentum after hopping. Depending on the exact pose of the probe with respect to the underlying terrain, varying degrees of slip may occur as the chassis rotates, thus inducing varying degrees of angular-to-linear momentum transfer. Spinning probes are not a problem for the gravity estimation, as they are symmetric by design and use an omnidirectional reflecting surface or LED array to emit sufficient light in any orientation. Rotating probes will, however, eventually dissipate their rotational energy during subsequent collisions, which can act to either increase or decrease the rebound speed, depending on the geometry of the collision. Thus, future work should study the probabilistic nature of residual angular momentum after hopping on various surfaces and its effect on subsequent bounces.

5.2.3 Hopping Operations

Now that we have derived the speed requirements for hopping and discussed several mechanisms that may be used to achieve these speeds, the final question is “when should they hop?” The seemingly obvious answer is “as soon as they come to rest on the surface.” Indeed that is one potential strategy would maximize the flight time of the probes, but it would also require additional sensors to determine whether or not the probe is landed or not, such as an IMU or contact sensors, which may be unreliable.

A predefined schedule for executing hops is an attractive option due to its simplicity and robustness; the probes need only trigger hops at certain predefined timestamps. This strategy requires no sensors and has the advantage of knowing precisely *when* the probes execute hops, allowing for spatial correlation of observed trajectories with the shape model of the body, improving gravity estimation (this is left for future work). Furthermore, the

probes could either share a single predefined hop schedule, or they could each have unique schedules, allowing for probe identification by noting the observed timestamps of each hop. However, a scheduled approach may inadvertently attempt to execute hops while the probes are aloft, wasting energy. Thus, a statistical analysis of how long each hop (and subsequent bounces) take, similar to the Monte Carlo study discussed in Sect. 5.2.1, should be conducted to select an appropriate hop frequency. If the host mission can afford longer observational campaigns, then lower hop frequencies may be acceptable. In fact, there may actually be value in having the probes resting on the surface for some period of time since they provide body-fixed landmarks that may improve the spacecraft’s vision-based navigation, and thus improve gravity estimation. Future work should investigate the secondary use of probes as artificial surface landmarks.

Finally, If the probes were able to receive commands from the mother spacecraft, hops may be requested directly. The advantages of this approach would be that (1) the commands could be individually addressed, allowing for simple probe identification, and (2) only probes on the visible side of the body would hop. However, this strategy could also present a risk that the probes settle in a local concavity such that they are never within line-of-sight of the spacecraft to receive commands. Thus, this strategy should be coupled with an auxiliary “time-out” trigger so that probes don’t wait indefinitely in these cases.

In summary, (1) speed requirements for suborbital hopping on small bodies are modest and can be derived analytically with simple equations or probabilistically via Monte Carlo simulations, (2) a number of simple mechanisms exist that could achieve random hopping at these speeds, and (3) the operational logic for triggering hops can be realized as a simple predefined schedule.

5.3 Probe Visibility

The second key requirement for the probes is that they be “visible” to a camera onboard the mother spacecraft. In other words, they need to be bright enough to stand out as a detectable point source against the background, which may either be dark sky (if the probes are beyond the limb of the body) or the small body itself, illuminated by the sun. Two methods were considered for achieving this: (1) passive reflection of sunlight and (2) light emission. A third strategy using retroreflective materials was initially considered, requiring a flash lamp on the mother spacecraft. This method was successfully employed on Hayabusa2 to track retroreflective “target markers” deployed to the surface during the terminal descent [54], but it’s performance degrades as $1/r^4$ and would require an unreasonably powerful flash lamp to “outcompete” the solar illumination at the $> 1 \text{ km}$ distances of an orbiting mother spacecraft in this concept. Thus, this section focuses on the more realistic strategies of sunlight reflection and light emission.

5.3.1 Light Reflection

The simplest strategy for achieving the brightness requirement for optical tracking is by passive reflection of sunlight. Indeed this is the most common paradigm in various tracking applications such as orbital debris tracking from ground-based telescopes [55] and of course in tracking solar system objects, including planets and small bodies. There have even been other concepts proposed to perform flyby gravity science of small bodies with deployed probes

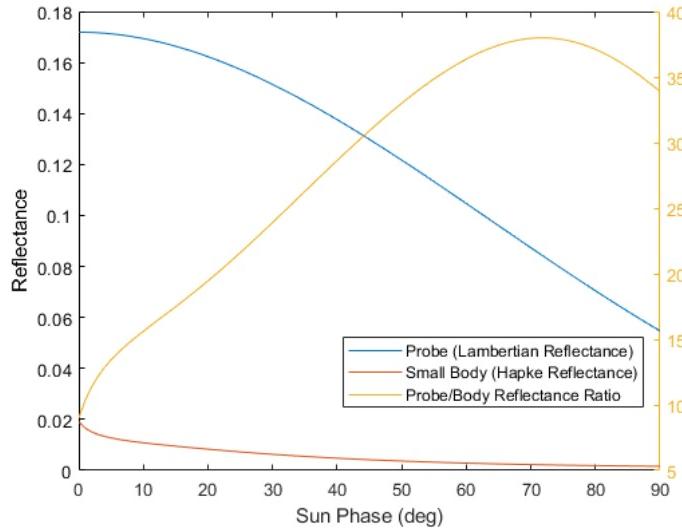


Figure 20: The brightness phase functions of a diffuse (Lambertian) white spherical probe and a small asteroid with a Hapke reflectance model. Their ratio is also plotted to show the relative brightness of the probe with respect to the asteroid background.

that reflect sunlight [56], [57]. However, in nearly all cases, reflecting objects are assumed to be detected against a dark sky background, which presents very little noise in comparison to an illuminated small body surface. Thus, Gravity Poppers that rely on sunlight reflection to be detected in front of the small body will need to be comparatively larger in order to achieve sufficient SNR for detection.

Before deriving the size requirements for reflecting probes, we must first consider their shape and optical reflectance properties. Many geometries could be considered, but since the probes are not attitude controlled and could therefore be in any arbitrary spin state, a *spherical* geometry is ideal for the directional uniformity of reflected light. As for material properties, the goal is to select an optical coating (or combination of coatings) whose “Bidirectional reflectance distribution function” or (BRDF), when applied to a spherical probe, is detectable at the widest range of solar phase angles, allowing for a wider range of orbit / station keeping geometries for the host spacecraft. A variety of high-albedo reflecting materials were considered including specular (“shiny”), diffuse (“flat”), and retroreflective materials, each of which has different characteristic phase functions when applied to a sphere. We found that the broad reflectance spectrum of a diffuse white sphere maximizes the solar phase angles at which the probe can be detected against the body.

Figure 20 shows the relative BRDF of a diffuse white sphere with an albedo of 0.9 and a small asteroid with albedo of 0.05 and a Hapke BRDF model [58]. The relative brightness ratio is also shown. Interestingly, although the diffuse probes are dimmer at larger solar phase angles, the brightness of the body decays even faster, thus leading to a higher brightness ratio at large solar phase angles. However, even at a 0° solar phase angle, at which the body exhibits a brightness surge due to the well-known opposition backscattering effect of regolith surfaces, the probe is still 10 times brighter. A small amount of retroreflective materials (e.g. those used for traffic signs) could be added to the probes to improve their relative

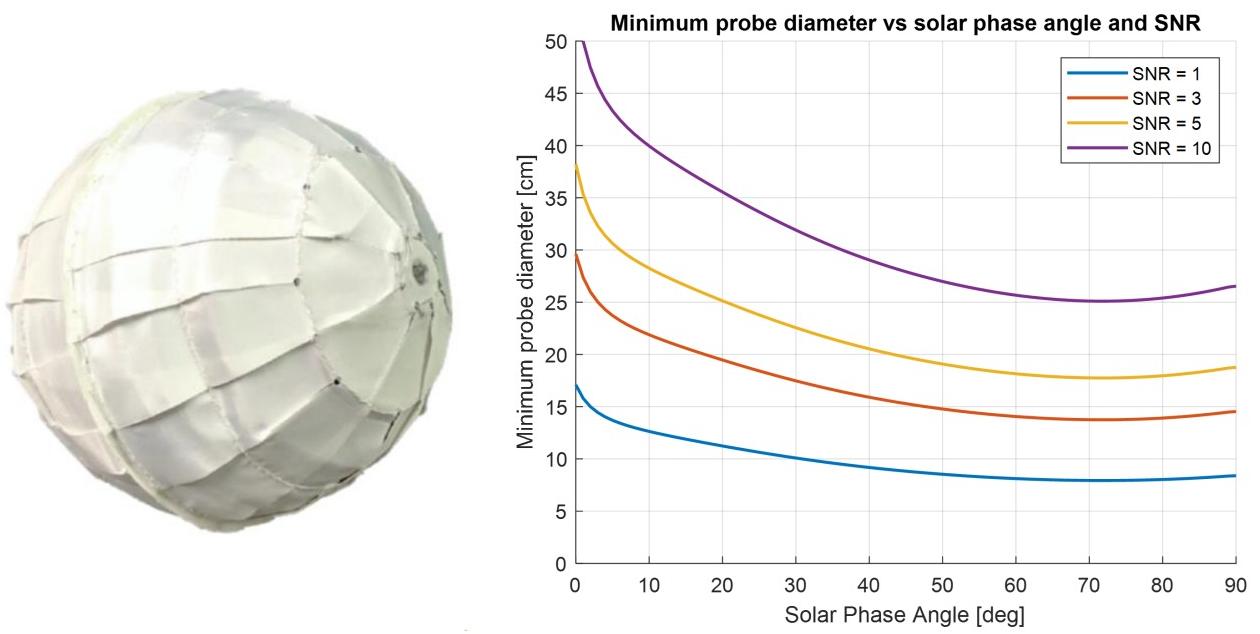


Figure 21: **Left:** Passive reflecting spherical probe concept with a with diffuse coating (credit: [56]) **Right:** The minimum diameter of a diffuse white spherical probe to be detected against a small body background (Bennu in this case) as a function of solar phase angle and SNR.

brightness near solar opposition, but at the cost of brightness at larger solar phase angles. Specular materials may enhance detectability at very high solar phase angles ($> 90^\circ$), but the uniformity of specular reflection is highly sensitive to irregularities in their spherical shape (i.e. the “disco ball effect”). Thus, a purely diffuse white sphere was selected as the baseline option (Fig. 21-Left).

Now given the brightness ratio of the diffuse spherical probes as a function of solar phase angle, $B(\phi)$, and assuming an imaging geometry whereby the frame half-width is k_{im} times the body radius, R_B , the minimum diameter of a probe is given by

$$D_p(\phi) = \frac{4k_{im}R_B}{n_{px}} \sqrt{\frac{S_{min}\eta}{\pi(B(\phi) - 1)}}, \quad (6)$$

where n_{px} is the pixel width of the sensor, S_{min} is the minimum SNR for detection, $\eta = 1/0.68$ is the factor accounting for defocused optics. The first key insight from this equation is that the minimum probe diameter scales linearly with the physical span of one pixel (i.e. R_B/n_{px}), and thus proportional to the size of the body. For our Bennu case study, where $R_B = 250$ m, $k_{im} = 1/0.7$, and $n_{px} = 2048$, Fig. 21-Right shows D_p as a function of solar phase angle and S_{min} .

This analysis suggests that for passively reflecting probes to be detectable in front of the lit small body, they may need to be quite large. For a C-type asteroid with mean albedo of 0.05, $D_p \approx R_B/2500$, but for brighter S-type asteroids with an albedo of 0.25, $D_p \approx R_B/1000$. Thus, passive reflection may be an unwieldy solution for large and/or bright bodies. The high-albedo diffuse coatings of the probes may also be subject to degradation during collisions with a dusty environment. However, this analysis assumed a simple thresholding filter for

probe detection. More advanced background subtraction methods may be able to enable detection of fainter, and thus smaller, probes.

5.3.2 Light Emission

To overcome the challenges associated with reflected sunlight for probe detection, *emitted* light through LED strobining may be considered. LED strobing has many advantages over passive sunlight reflection. First, it is insensitive to solar illumination angles and can even allow for detection when the probes are in eclipse/shadow. Second, while a strobing probe still has to compete with the background brightness of the small body, an appropriate selection of narrow spectrum LEDs and a *matched band-pass filter* on the camera can allow for up to 100x reduction in background noise, dramatically improving SNR (see Fig. 22). Third, LED's are less susceptible to dust degradation than reflective optical coatings. Fourth, in contrast to reflecting probes which are always “on,” the ability take calibration images when the probes are off may allow for more effective background subtraction. Finally, LEDs offer the potential to encode unique identifiers through frequency modulation, allowing for simpler probe identification.

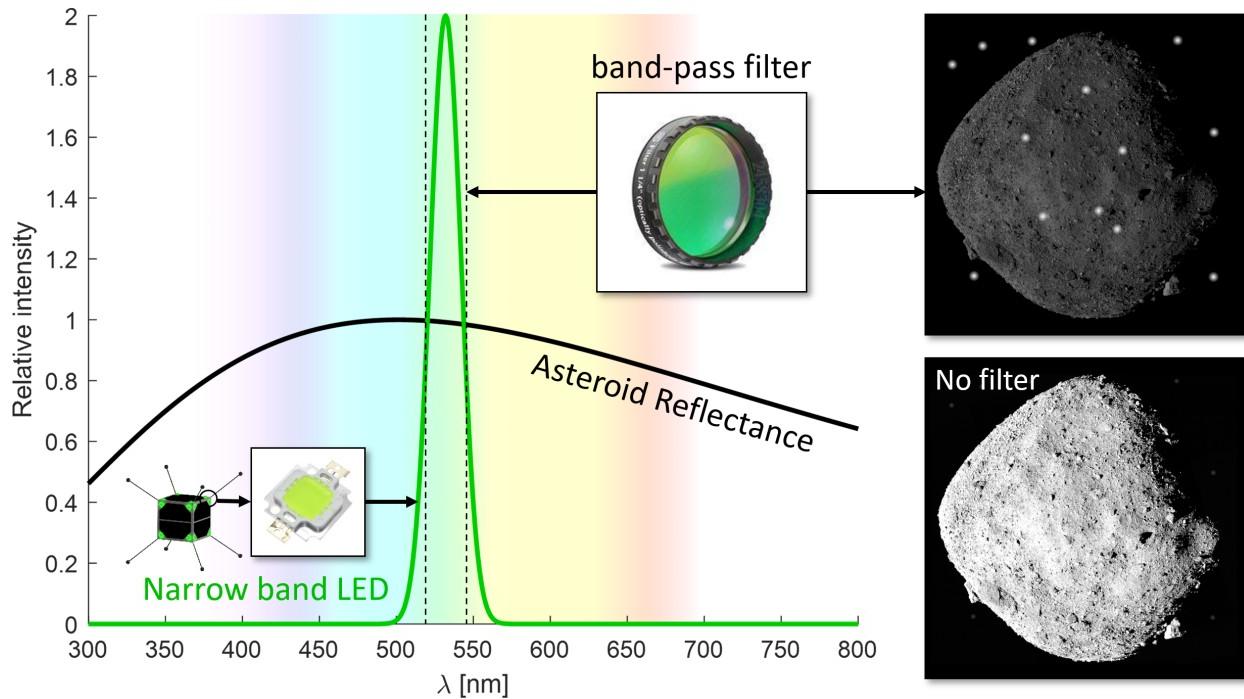


Figure 22: An illustration of the SNR improvement that can be provided by a narrow-spectrum LED with a matched band-pass filter. Compared to a broad spectrum light source and sensor, an narrow band-pass filter can allow for detection of specific LEDs that are 50 to 100 times dimmer!

The principle of using strobbed LEDs for optical tracking is not a new concept. It is the basic principle of many motion capture systems and has even been demonstrated in space. The Japanese LEO CubeSat, FITSAT-1, was equipped with an array of strobing LEDs that could be detected at ranges over 400 km using modest telescopes [59]. While the scales and background noise characteristics are quite different than the Gravity Poppers concept, this project validated the technique of detecting weakly emitted light over vast distances through

an appropriate selection of optical filters and exposure times.

The technique of strobing LEDs does present a few challenges, however. The obvious burden is their *power* requirements, which is discussed in detail in Sect. 5.4. Second, because the probes are not attitude controlled, an omnidirectional array is required, which is less efficient than a pointable probe like FITSAT-1, which only required LEDs on one side. Third, the center of an LED’s emission spectra can vary with temperature, thus limiting the minimum band-pass filter window than can be used to reduce background noise. Finally, the synchronization of pre-programmed LED flashes with the camera shutter can be challenging, especially for simple clock oscillators that drift due to thermal variations on the order of 1 ppm. However significant clock drift may be accommodated in a number of ways, such as (1) using a very high frame rate sensor (2) estimating the clock drift of each probe and adjusting the image timing accordingly, or (3) by avoiding a clock altogether by querying each strobe with a simple modulated radio chirp from the mother spacecraft.

In summary, two methods for achieving the brightness requirements for optical tracking were considered: sunlight reflection and LED strobining. Passive reflection has the advantage of not requiring any power but was determined to require a large surface area (in proportion to the body size) which may be difficult to stow in a small volume. Probes with strobining LEDs have many advantages over reflecting probes, and, as discussed in the next section, have manageable power requirements. Thus, we recommend the use of strobining LEDs for Gravity Poppers that can be used in the broadest range of environments.

5.4 Power System

The power subsystem may be one of the most critical considerations for the probes, as it is a primary driver for their minimum size (i.e. to accommodate a battery and possible solar cells). Moreover, power requirements for both LED strobining and hopping are highly dependent on certain properties of the small body environment. Thus, a study of the power system may reveal key scaling laws that relate the minimum probe size to key environmental parameters.

5.4.1 Hopping Power

The average power requirements for continuous hopping can be described as the energy required for executing each hop multiplied by the hop frequency. For a first order analysis, we set the hop energy to half of the escape energy of the body, and the time between hops as the orbital period for a hop with a semi-major axis equal to the body radius. Accordingly, the average hopping power can be described as:

$$E_{hop} = \frac{4\pi G m_p \rho_B R_B^2}{3\varepsilon_{hop}}, \quad T_{hop} = \frac{\pi k_b}{\sqrt{G\pi\rho_B/3}}, \quad \rightarrow \quad \bar{P}_{hop} = \frac{E_{hop}}{T_{hop}} = \frac{4m_p R_B^2 \sqrt{\pi(G\rho_B/3)^3}}{k_b \varepsilon_{hop}}, \quad (7)$$

where G is the gravitational constant, R_B and ρ_B are the radius and density of the small body, m_p is the probe’s mass, nominally set to 200g, k_b is a time multiplication factor to allow the probe to come to rest on the surface, nominally set to 2, and ε_{hop} is the “hop efficiency” (i.e. from electrons to kinetic energy), nominally set to 0.01 (consistent with energy efficiencies of internal momentum devices, e.g. [36]).

As expected, hopping power increases with the size and density of the small body, suggesting that practical limits on body size may be driven by power constraints. However, for the body size range of interest for this concept, hops are quite infrequent and require very little energy. Thus, even for a large asteroid like Eros ($R_B = 8.4\text{km}$, $\rho_B = 2700\text{kg/m}^3$), the average power is only **30 mW** and for a body as small as Bennu ($R_B = 250\text{m}$, $\rho_B = 1200\text{kg/m}^3$), it is a minuscule **10 μW**. Thus, the energy requirements for hopping may only be a concern for very large asteroids and if the hopping mechanism is particularly inefficient (say, $\varepsilon_{hop} < 0.001$).

5.4.2 Strobing Power

For a detection architecture requiring the probes to strobe an LED, we must account for the power required to do so in the energy budget. As discussed in Sect. 5.3, two regimes are possible for optical detection: (1) when the probe is in front of a dark sky background and (2) when the probe is in front of the lit small body. The former requires far less power as the background noise is negligible (within the narrow band-pass filter range). It can be shown that the energy required for a detectable LED flash against a dark sky background is:

$$E_{LED} = \frac{8e_\lambda k_{im}^2 R_B^2 S_{min}^2}{Q_e \varepsilon_{LED} D_{cam}^2 \tan^2(\alpha_{cam}/2)} \left[1 + \sqrt{1 + 4\pi \left(\frac{N_r \sigma_{px}}{S_{min}} \right)^2} \right], \quad (8)$$

with parameters and (nominal values) as follows: $e_\lambda = hc/\lambda_{LED}$ is the LED photon energy (**3.7e-19 J** for a green LED), k_{im} is the multiplicative margin on the body radius (R_B) for an image frame (**2**), S_{min} is the minimum SNR for detection (**5**), Q_e is the sensor quantum efficiency at λ_{LED} (**0.5**), ε_{LED} is the LED radiant efficiency (**0.3**), D_{cam} is the aperture diameter of the camera (**0.06 m**), α_{cam} is the camera field of view (**20°**), N_r is the read noise (**10 e⁻**), and σ_{px} is the PSF radius in # pixels (**3**). This formulation assumes that the exposure time is sufficiently short such that dark current is not a significant source of noise.

With these values and for body radius of 250 m (e.g. Bennu), $E_{LED} \approx 30 \mu\text{J}$, which could be realized, for example, by a 50 mW pulse for 1 ms. For Eros ($R_B = 8.4\text{km}$), $E_{LED} \approx 50 \text{mJ}$, or a 50 W pulse for 1 ms. For a 0.1 Hz pulse frequency, for example, this is still a very small average power of 5 mW. It is not too surprising that the minimum power requirements for detection against a dark sky are so low; after all, OSIRIS-REx could detect cm-scale, low-albedo rocks with sufficient exposure times.

However, long exposures to observe dim probes around the body can create other issues, such as the “bloom” effect of the body, whereby overexposed pixels bleed into neighboring pixels, thus requiring more complex background-subtraction algorithms. Thus, it is preferred to operate in the second regime, in which the probes are detectable without any overexposure, even when in front of the lit body. Although more sophisticated background subtraction methods may be used to “remove” the background image of the small body, for simplicity, we will assume a basic thresholding filter whereby the pixels lit by the probes must be brighter than the body image by a factor of S_{min} . Thus, it can be shown that a probe flash must have a minimum power of,

$$P_{flash} = \frac{16\pi S_{min} q_s k_{im}^2 R_B^2 w f_\phi f_{bp}}{a_B^2 n_{px}^2 \varepsilon_{LED}} \quad (9)$$

with additional parameters and (nominal values) as follows: q_s is the solar constant (1361 W/m^2), a_B is the distance of the small body from the sun (1.2 AU for NEAs), w is the geometric albedo of the body (0.05 for C-type asteroids), f_ϕ is the reflected fraction of light at solar phase angle ϕ (0.2 sr^{-1} at $\phi = 20^\circ$ for mean C-type Hapke reflectance model), f_{bp} is the fraction of reflected sunlight transmitted through the pass-band filter (0.025 for $520 - 540 \text{ nm}$ filter), and n_{px} is the number of pixels in the image span (2048 for a standard 4 MP sensor). For a Bennu-sized body, this suggest a minimum strobe power requirement of **5 W**. High-power COTS LEDs are readily available with strobe power ratings in excess of 10 W/cm^2 , suggesting that P_{flash} may be as high as 1 kW for a 5 cm probe, before requiring more surface area.

For larger and/or brighter bodies where the LED strobe power may be a concern, several measures may be taken to reduce these requirements, as suggested by (9). First, increasing the number of pixels of the camera (n_{px}) reduces the brightness requirement, as this reduces the physical span of reflected light from the body per pixel. Second, more narrow-band, thermally stable LEDs may be considered, allowing for a narrower band-pass filter on the camera, and thus reducing f_{bp} . Third, operationally, the tracking may be restricted to larger solar phase angles (reducing f_ϕ) so that the body is dimmer. Fourth, this analysis assumed an omnidirectional array of LEDs; however if the probe is equipped with a directional sensor (e.g. an IR photodiode array), it may selectively strobe only the spacecraft-facing LEDs. Finally, and perhaps most importantly, more advanced background subtraction techniques may be employed (e.g. by taking “calibration frames” just before and after each strobe) so that the probes need not be brighter than the body, only brighter than the shot noise of the background-subtracted image.

Finally, to compute the average power of the strobing LED, we need to multiply (9) by the LED on-time, or accordingly, the camera exposure time (t_{exp}), and the strobe frequency (f_{img}), nominally set to 0.1 Hz as per the 10-second imaging frequency assumed in Sect. 3. When setting the camera exposure to maximize dynamic range, the average LED power can be expressed as:

$$P_{LED} = P_{flash} t_{exp} f_{img} = \frac{16W e_\lambda k_{im}^2 R_B^2 S_{min}}{Q_e \varepsilon_{LED} D_{cam}^2 \tan^2(\alpha_{cam}/2)} f_{img} \quad (10)$$

where W is the full well capacity of the sensor, nominally set to $10,000$. For the relevant values in our Bennu case study, $P_{LED} \approx 1 \text{ mW}$, a fairly modest average power draw despite the high-power bursts.

5.4.3 Hotel Power

In addition to the power requirements for the key functions of hopping and strobing, a minimal set of other electronic components are required for control. Namely, a microcontroller is required for executing the simple logic / state machine for strobing and hopping and an external stable oscillator may be required for minimizing clock drift. Small, low-power microcontrollers (e.g. “PIC 16F”) are readily available with power consumptions under **1 mW**. Oscillator types vary considerably, but a notional choice for such a small probe might be a “temperature compensated crystal oscillator” (TCXO), many of which have power draws less than **10 mW**. Finally, although not baselined in this Phase I study, an ultra-low-power

radio receiver may be accommodated on the probe for simple commanding, which generally consume a few 10's of mW.

In summary, the average power required for hopping and strobing are described by equations (7) and (10) respectively, which for our Bennu case study, correspond to powers of **10 μ W** and **1 mW**. “Hotel” loads for a low-power microcontroller and oscillator account for another **10 mW**. However due to the first-order nature of this analysis (and the possibility to add a radio receiver), we can add an order of magnitude margin to this, yielding a maximum expected value (MEV) of **100 mW** mean system power draw.

5.4.4 Power System Architecture

Two options exist for the probe’s power supply: (1) a primary, non-rechargeable battery, or (2) a secondary battery with solar cells for recharge. Given an average power demand of 100 mW, a primary battery with energy density of 500 Wh/kg (e.g. Li-(CF)_x) and a mass of only 5 g would be sufficient to power the probe for the nominal 24-hr mission simulated in Sect. 3. However, a 5 cm probe could easily accommodate a 50 g battery, which would provide sufficient power for over 250 hours of operation. Secondly, should the desired lifetime be longer than the capacity of a primary battery or the probe’s power demands grow, a secondary battery and solar cells may be considered for indefinite operation. At 1 AU, practical solar arrays can provide about 250 W/m². Therefore, a 5 cm probe can generate about 600 mW of continuous power. Even conservatively assuming that the probe is in eclipse/shadow for half of the time, this is still sufficient to power the probe indefinitely. A small secondary battery could then be sized for the longest expected period of shadow/eclipse (e.g. 1 rotation period); a 5 g coin-cell battery would be more than enough. Capacitors are also required to provide burst power for LED strobing.

An initial consideration for achieving hopping without any electronics was to extract energy from the changing thermal environment (i.e. as the probes pass in and out of eclipse) in the form of mechanical strain (e.g. by “charging” a shape-memory alloy spring). We determined that occurrences of sun/shade transition events (required to charge and discharge such a device) were too sporadic and unreliable for randomly hopping probes. We concluded that such an architecture would lead to a high failure rate when the probes get “trapped” in permanently shadowed or lit regions with insufficient thermal variations. Given the very low power requirements of hopping and strobing, we determined that simple electrical power system is more favorable than thermal-mechanical system.

5.5 Thermal System

The thermal environment of small bodies can often be very harsh and volatile. Probes that hop in and out of eclipse and with varying degrees of IR exposure from the small body may experience dramatic temperature swings. However, the probes are far too small and power-limited to employ any active thermal management, and without attitude control they cannot leverage directional coatings either. Thus, the objective of the thermal design is to ensure the passive equilibrium temperatures of the probe remain within the operating range of all subsystems when passively radiating to the environment. A conservative range for common electronics would limit operations within -20 to 50°C; however, temperature-tolerant batteries, microcontrollers, motors, and LEDs exist that can operate from -40 to

150°C, albeit typically at reduced performance.

The general equation that governs the equilibrium temperature of a probe (approximated as a sphere) in space is

$$T_{probe}[K] = \left[\frac{B_{sun}q_s\alpha/a_B^2 - Q_p/(\pi r_p^2)}{4\sigma\varepsilon} + F_B T_B^4 \right]^{1/4}, \quad (11)$$

where α and ε are the mean solar absorptivity and IR emissivity of the probe's surface, B_{sun} is a binary variable for if probe is in shadow (0) or in sun (1), Q_p is the internal heat generation, r_p is the probe's mean radius, $\sigma = 5.67e-8$ is the Stefan-Boltzmann constant, F_B is the view factor of the small body from the probe, and T_B [K] is the mean effective temperature of the small body surface that is visible to the probe.

The first key observation is that the distance of the small body from the sun, a_B , is a driving variable, as it directly affects the solar incidence on the probes and the surface temperatures of the body, T_B (which is also affected by the albedo and thermal inertia). Depending on the body, the range of surface temperatures, ΔT_B , can vary from 70°C for faster rotators with high thermal inertia to over 200°C for bodies with low thermal inertia or high obliquity. For Bennu, T_B ranges from about 200 K to 350 K. The second key observation is that the probe's surface coatings (α, ε), which are the main design variables for passive thermal management, have a large impact on the relative effect of solar absorption.

The first step in thermal design is to understand *where* the probes are spending their time. Of course, randomly hopping probes could go anywhere, from resting on the hot surface to a distant eclipsing orbit. However, it is helpful to define a mean altitude at which the probes are at *on average*, so that a mean view factor \bar{F}_B of the body can be used to estimate equilibrium temperatures. According to our dynamic simulations in Sect. 5.2.1, mean altitudes range from about $0.5R_B$ to R_B , corresponding to a mean view factor of about $\bar{F}_B = 0.2$. Now we can use this mean view factor and a global mean temperature (of say, $\bar{T}_B = 270$ K) to select surface coatings (α/ε) to achieve a desired equilibrium probe temperature, say $\bar{T}_{probe} = 20$ °C. For this example, $\alpha/\varepsilon = 1.25$, which is well within the range of selective absorbing coatings.

Now that we've addressed the mean case to select an α/ε for nominal operations, we can compute the worst-case-hot (WCH) and worst-case-cold (WCC) temperatures. WCH occurs when the probe is at rest on the hottest part of the surface ($F_B = 0.5$) and in direct sunlight. For the parameters defined above, and a max surface temperature of 350 K, this corresponds to a probe temperature of $T_{probe, max} = 70$ °C. WCC occurs in a distant eclipse, where $B_{sun} = 0$ and $F_B \rightarrow 0$. In this case, a steady state analysis suggests a probe power of about $Q_p = 1$ W is required to keep $T_{probe, min} > -40$ °C. However, eclipses are generally relatively brief and may be survived simply via the thermal mass of the probe, or optionally, using a small amount of phase-change material. For more distant bodies such as main belt asteroids where the cold case may be more driving, insulation can be considered to keep the probes warm.

In summary, a first-order thermal analysis suggests that the probes may be designed to operate passively within their temperature limits with an appropriate selection of thermal coatings. Future work should consider a statistical analysis of probe transient temperatures as they hop around the small body, including shadows cast by the small body and a thermophysical model of the spatial/temporal temperature distribution. This will help to assess

more realistic temperature ranges for the probes and further inform design measures such as selective coatings and possibly phase-change material. Operational logic may also be investigated to help regulate temperatures, for example, by choosing when to hop based on the need to cool down or heat up.

5.6 Notional Probe Concepts

Based on the requirements derived above for key subsystems (hopping, visibility, power, and thermal), two notional probe architectures were conceived (see Fig. 23), one which uses passive sunlight reflection and one which uses an active LED strobe to be detected by the tracking camera.

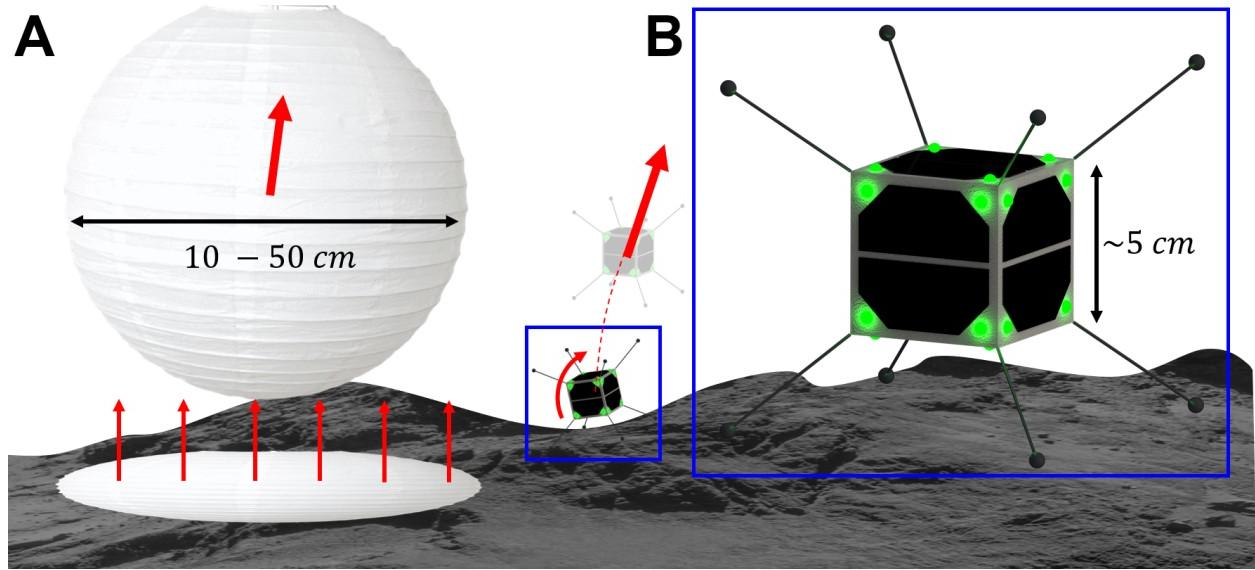


Figure 23: Two Gravity Popper probe design architectures derived in the phase I study. **(A)** A diffuse white sphere reflects sunlight to be detected by the spacecraft camera and uses an internal linear spring mechanism to hop. **(B)** A “femto-Sat” probe design which uses an omnidirectional array of strobing LED’s to be detected by the spacecraft camera and an internal reaction wheel for hopping.

The reflecting spherical probe uses a simple internal linear mechanism to wind a spring, collapsing it to a disk, and then releasing it abruptly. The shape expansion hopping mechanism allows the probe to “self-right” automatically, enabling it to hop on even the most complex terrains. However, while this probe is simple and robust, for all of the reasons discussed in Sect. 5.3, sunlight reflection requires the probes to be quite large, making them more massive/voluminous to accommodate—not ideal for a minimalistic mission concept. Thus, we chose to focus more attention on the second concept in Fig. 23-B and further detailed in Fig. 24, which uses strobbed LEDs for detection.

The second probe geometry is a 5cm cube, which has sufficient volume to house all key subsystems, is an efficient shape for packaging (both of internal components and of the probes themselves within a deployment canister), and is a favorable symmetric geometry for predictable photon forces in any rotation state. Each external face of the probe has several LEDs with a wide-angle emission, though the exact number and geometric arrangement of these LEDs is left for future design work. The faces also have small solar cells to extend the

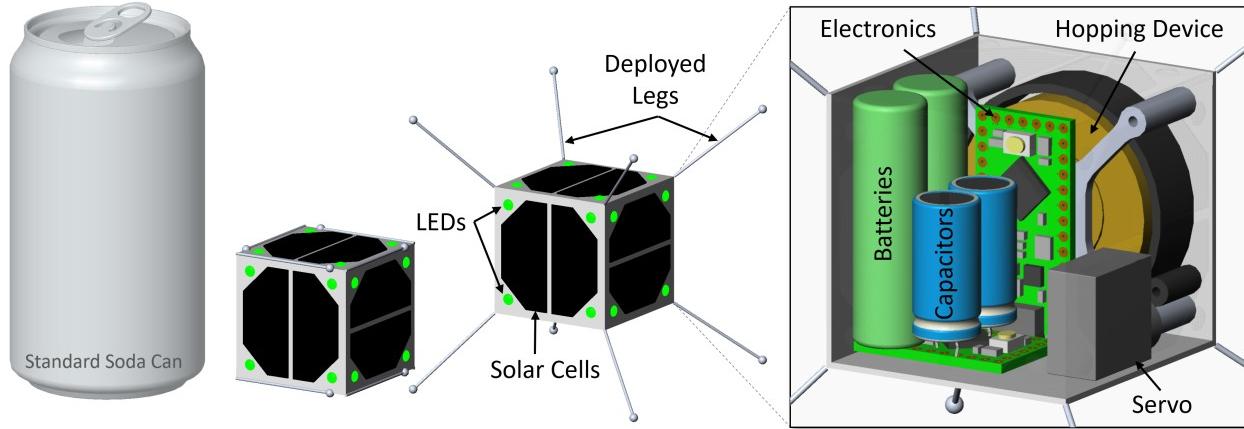


Figure 24: Notional design of a hopping probe for the Gravity Poppers mission concept, with a soda can shown for scale. The first image shows the probe in its stowed configuration. The second image shows the deployed configuration with eight radial landing legs on each corner. The third image shows a cutaway revealing the basic components within the probe: a flywheel hopping device, batteries and capacitors for energy storage, and simple electronics for power distribution and control.

lifetime of the probes beyond the capacity of the battery, which provide about 600 mW (at 1 AU) in any rotation state.

Externally, the probe also has eight thin “legs” (one on each corner) which are passively deployed upon ejection from a deployment canister, and which elastically bend during surface collisions in order to minimize energy dissipation and encourage longer bouncing. The cubic arrangement of these landing legs shown in Fig. 24 is highly notional and could be modified to, for example, an asymmetric geometry that preferentially biases the probe to settle in a particular configuration. Their 5 cm length is also highly notional and could be extended to provide greater elasticity or stand-off distance from the surface, or shortened to improve hop efficiency and packaging simplicity. Future work should consider these trades in the context of a higher-fidelity mobility system design and analysis.

Internally, the probe has a single hopping actuator consisting of a central, low-power brushless motor embedded within a large flywheel, and a small servo motor (or solenoid) to cinch a circumferential band to the flywheel for braking. The power system consists of two small secondary lithium batteries, with a total capacity of about 6 Wh, and auxiliary capacitor(s) for providing burst power for the LED strobes. As per the power analysis in Sect. 5.4, this secondary battery can power the probe for approximately 100 hours on a single charge, but also has the capability to recharge via solar cells for an extended mission. Finally, a simple electronics board consists of power distribution, a microcontroller, and a motor driver. Overall the system has a mass of about 170 g—about the same as a modern smart phone—and a stowed volume of only 125 mL, allowing eight probes to be stowed within a standard 10 x 10 x 10 cm (1U) CubeSat volume.

Finally, this minimalist design assumed no communication capabilities with the mother spacecraft (e.g. to receive hop or strobe commands). However, miniaturized low-power radio technology is constantly improving and may be a consideration in future designs. The deployed landing legs may even be able to serve a dual purpose as an omnidirectional dipole antenna.

5.7 Probe Deployment System

While not a primary focus of the Phase I study, we designed a notional probe deployment canister and to illustrate how the probes (and accompanying tracking system) might be accommodated on a host space spacecraft. A simple CAD design of the “Gravity Poppers” instrument concept is shown in Fig. 25. The probe canister has a 3U (i.e. 10 x 10 x 30 cm) volume and accommodates 24 probes. It has a simple linear spring mechanism that provides a preset ΔV to the whole batch of probes during a single deployment event. This mechanism is very similar to many heritage CubeSat deployment systems, which have only a few percent error on the ejection velocity—sufficient for robust gravitational capture of the swarm at an appropriate deployment altitude. Figure 26 illustrates what a deployment may look like to the surface of Asteroid Bennu from a fixed hovering altitude.

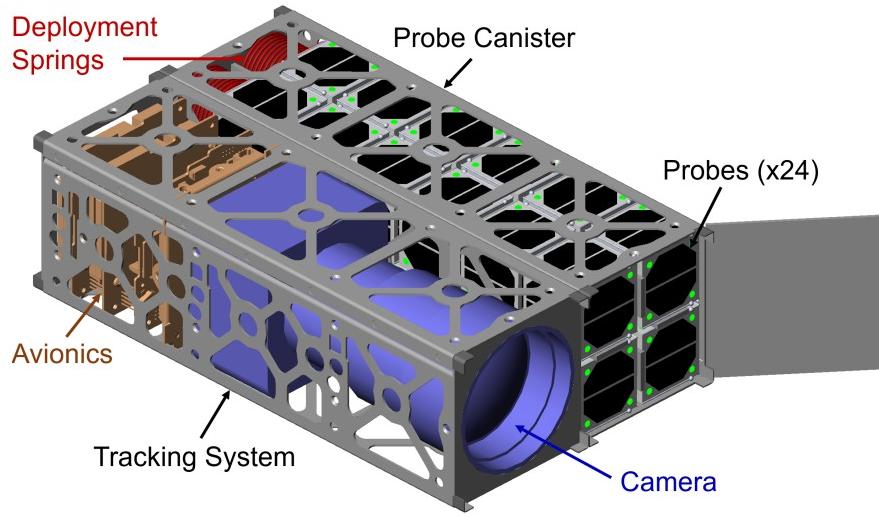


Figure 25: Notional CAD design of the Gravity Poppers instrument concept, conceived as a 6U form factor: 3U dedicated to the probe deployment canister and 3U for the tracking system (camera and avionics for image processing).

Next to the probe canister is a notional design of the probe tracking system, which consists of a single, astrometric-quality camera (this 2U camera model is based on the Asteria CubeSat [60]), and a few avionics boards for power distribution, image processing, and data handling. This design assumes that the host spacecraft provides ($< 10\text{ W}$) power and is equipped with its own attitude determination system. However, this design could be easily augmented with one or more star trackers to provide a fully self-sufficient orbit and attitude determination solution. Overall, this Gravity Poppers instrument concept is only 7 kg (4 kg for the probes and 3 kg for the structures and tracking system), which is small enough to be accommodated on a number of spacecraft that may be visiting small bodies.

6 Conclusions

In this effort we investigated a novel mission architecture for high-resolution gravity recovery of small solar system bodies—a measurement that is notoriously difficult for small bodies using traditional gravity science methods. Our study demonstrated that a swarm of small hopping probes can in fact be used as effective “gravity beacons.” Initial simulation studies

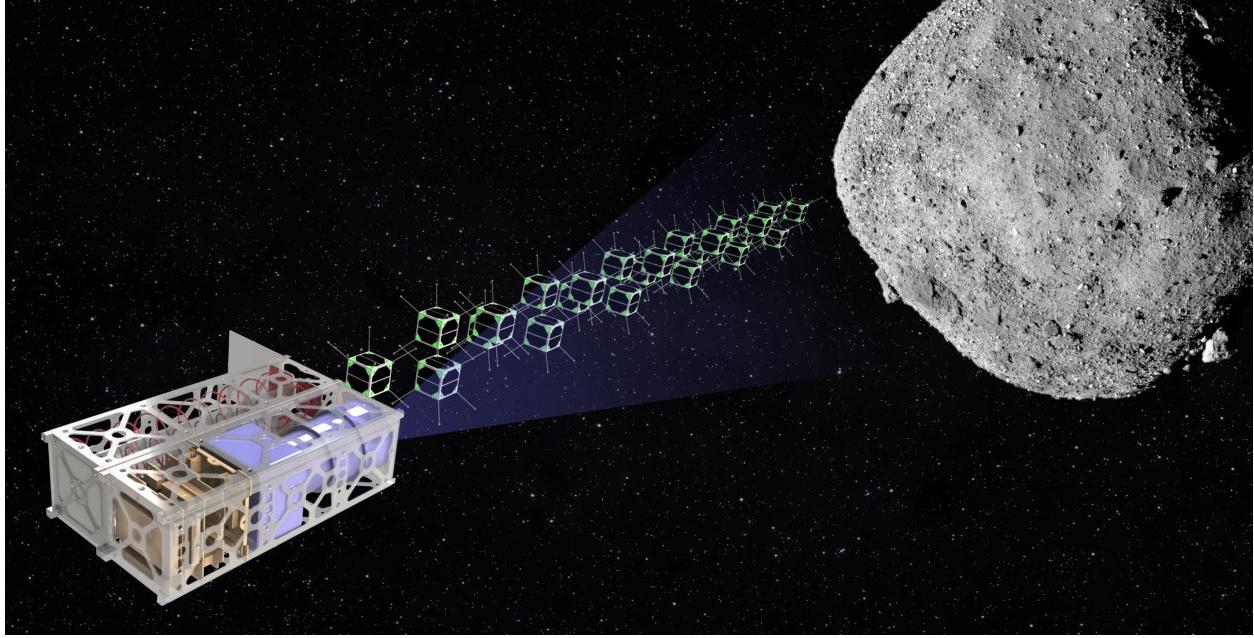


Figure 26: Artist concept of Gravity Poppers being deployed to the surface of Asteroid Bennu.

suggest the ability to recover gravity fields beyond degree-and-order 100 for a small, 500 m body, orders of magnitude better than existing methods, and sufficient to address a number of interesting science questions relating to the interior structure and composition of small bodies.

In addition to the technical feasibility of the concept, the minimal hardware requirements revealed through our systems engineering analysis suggest that this concept could be realized as a small *instrument* payload (<10 kg), which could be hosted on a variety of spacecraft visiting small bodies, rather than a dedicated mission. Thus, this effort revealed that Gravity Poppers is not only enabling for small body gravity science, but also that there is a high potential for mission infusion.

Future studies should focus on the key feasibility and maturation aspects identified during Phase I, which can be grouped into the same three categories:

- **Swarm Gravity Estimation:** The backbone of this concept is a thorough understanding of how a tracked swarm of hopping probes translates into gravity field observability. This Phase I effort made many assumptions and simplifications in our simulation studies that should be relaxed in the future. In particular, future gravity estimation analysis should (1) consider bodies of arbitrary shape, (2) develop gravity inversion and mass estimation algorithms based on swarm tracking data, (3) refine the treatment of uncertainties associated with pointing, tracking, and perturbation forces, (4) extend the sensitivity analyses to a wider range of tracking, dynamics, and environmental parameters, and (5) demonstrate the ability to recover ground-truth internal mass distributions through blind experiments, validating the science value of the concept.
- **Tracking System:** The tracking system is the key interface between probe dynamics

and gravity estimation. This Phase I effort conducted a basic trade study on tracking sensors and paid particular attention to an optical tracking system which showed promising potential for its simplicity. However future studies should (1) revisit the trade study of tracking sensors in the context of a refined gravity observability analysis, (2) develop higher fidelity models of tracking uncertainties, (3) demonstrate automated data/image processing algorithms on high-quality synthetic imagery, and (4) Validate optical tracking in a hardware test bed or field campaign.

- **Probe Design and Systems Engineering:** The hopping probes are at the heart of this instrument concept and require significant maturation beyond this Phase I study. In particular, future studies should (1) prototype and test different hopping mechanisms in reduced-gravity test beds and on a variety of asteroid and cometary surface analogs, (2) mature the system design of the hopping probes, including the LEDs, power system, and avionics, (3) develop an integrated functional prototype demonstrating that small-size packaging can be achieved, and (4) conduct a comprehensive mission risk analysis to understand how the Gravity Poppers architecture may be made more robust.

7 Acknowledgements

The research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration (80NM0018D0004). The investigators would like to acknowledge the contributions of graduate student, Jacopo villa, for leading the technical study of gravity estimation. We would like to thank collaborators Steven Chesley and Zachary Manchester for their helpful guidance and suggestions throughout this project.

8 References

- [1] E. Asphaug, E. V. Ryan, and M. T. Zuber. Asteroid interiors. *Asteroids III*, 1:463–484, 2002.
- [2] D. J. Scheeres, D. Britt, B. Carry, and K. A. Holsapple. Asteroid interiors and morphology. *Asteroids IV*, 745766:745–766, 2015.
- [3] M. J. Sonter. The technical and economic feasibility of mining the near-earth asteroids. *Acta Astronautica*, 41(4-10):637–647, 1997.
- [4] National Research Council. Defending planet earth: Near-earth-object surveys and hazard mitigation strategies. Technical report, National Academy Press, 2010.
- [5] G. J. Consolmagno, D. T. Britt, and R. J. Macke. The significance of meteorite density and porosity. *Chemie der Erde-Geochemistry*, 68(1):1–29, 2008.
- [6] B. Carry. Density of asteroids. *Planetary and Space Science*, 73(1):98–118, 2012.
- [7] DJ Scheeres, AS French, P Tricarico, SR Chesley, Y Takahashi, D Farnocchia, JW McMahon, DN Brack, AB Davis, R-L Ballouz, et al. Heterogeneous mass distribution of the rubble-pile asteroid (101955) bennu. *Science advances*, 6(41):eabc3350, 2020.
- [8] W. Kofman, A. Herique, Y. Barbin, J.-P. Barriot, V. Ciarletti, S. Clifford, P. Edenhofer, C. Elachi, C. Eyraud, and J.-P. Goutail. Properties of the 67p/churyumov-gerasimenko interior revealed by CONSERT radar. *Science*, 349(6247):aab0639, 2015.
- [9] A. Safaeinili, S. Gulkis, M. D. Hofstadter, and R. L. Jordan. Probing the interior of asteroids and comets using radio reflection tomography. *Meteoritics & Planetary Science*, 37(12):1953–1963, 2002.
- [10] R. E. Grimm, D. E. Stillman, P. Sava, and D. Ittharat. Radio reflection imaging of asteroid and comet interiors ii: Results and recommendations. *Advances in Space Research*, 55(9):2166–2176, 2015.
- [11] Michael JS Belton, Peter Thomas, J Veverka, Peter Schultz, Michael F A'Hearn, Lori Feaga, Tony Farnham, Olivier Groussin, Jian-Yang Li, Casey Lisse, et al. The internal structure of jupiter family cometary nuclei from deep impact observations: The “talps” or “layered pile” model. *Icarus*, 187(1):332–344, 2007.
- [12] T. Saiki, H. Sawada, C. Okamoto, H. Yano, Y. Takagi, Y. Akahoshi, and M. Yoshikawa. Small carry-on impactor of hayabusa2 mission. *Acta Astronautica*, 84:227–236, 2013.
- [13] P. Michel, A. Cheng, M. Küppers, P. Pravec, J. Blum, M. Delbo, S. F. Green, P. Rosenblatt, K. Tsiganis, and J.-B. Vincent. Science case for the asteroid impact mission (AIM): a component of the asteroid impact & deflection assessment (AIDA) mission. *Advances in Space Research*, 57(12):2529–2547, 2016.
- [14] D. J. Scheeres, E. I. Asphaug, J. Colwell, R. Dissly, P. E. Geissler, L. A. McFadden, V. Petr, R. Reinert, and H. Yano. Asteroid surface science with pods. In *Lunar and Planetary Science Conference*, 2003.

- [15] W. F. Huebner. Seismological investigation of asteroid and comet interiors. *Mitigation of Hazardous Comets and Asteroids*, page 234, 2004.
- [16] O. Robert, P. Lognonne, D. J. Scheeres, N. Goujon, M. Le Feuvre, A. Izzet, C. Blitz, and L. Bowman. Seismology on a small body: expected results for the BASiX discovery mission proposal. In *AGU Fall Meeting Abstracts*, 2010.
- [17] J. Plescia, O. Barnouin, D. Richardson, N. Schmerr, D. Lawrence, B. Denevi, C. Ernst, and H. Yu. APEX asteroid probe experiment. In *Lunar and Planetary Science Conference*, 2017.
- [18] B. D. Tapley, S. Bettadpur, M. Watkins, and C. Reigber. The gravity recovery and climate experiment: Mission overview and early results. *Geophysical Research Letters*, 31(9), 2004.
- [19] T. L. Hoffman. GRAIL: gravity mapping the moon. In *IEEE Aerospace Conference*, 2009.
- [20] R. S. Park, A. S. Konopliv, B. G. Bills, N. Rambaux, J. C. Castillo-Rogez, C. A. Raymond, A. T. Vaughan, A. I. Ermakov, M. T. Zuber, and R. R. Fu. A partially differentiated interior for (1) ceres deduced from its gravity field and shape. *Nature*, 537(7621):515, 2016.
- [21] R. S. Park, A. S. Konopliv, S. W. Asmar, B. G. Bills, R. W. Gaskell, C. A. Raymond, D. E. Smith, M. J. Toplis, and M. T. Zuber. Gravity field expansion in ellipsoidal harmonic and polyhedral internal representations applied to vesta. *Icarus*, 240:118–132, 2014.
- [22] J. K. Miller, A. S. Konopliv, P. G. Antreasian, J. J. Bordi, S. Chesley, C. E. Helfrich, W. M. Owen, T. C. Wang, B. G. Williams, and D. K. Yeomans. Determination of shape, gravity, and rotational state of asteroid 433 eros. *Icarus*, 155(1):3–17, 2002.
- [23] L. Jorda, R. Gaskell, C. Capanna, S. Hviid, P. Lamy, J. Ďurech, G. Faury, O. Groussin, P. Gutiérrez, C. Jackman, S. J. Keihm, H. U. Keller, J. Knollenberg, E. Kührt, S. Marchi, S. Mottola, E. Palmer, F. P. Schloerb, H. Sierks, J.-B. Vincent, M. F. A'Hearn, C. Barbieri, R. Rodrigo, D. Koschny, H. Rickman, M. A. Barucci, J. L. Bertaux, I. Bertini, G. Cremonese, V. Da Deppo, B. Davidsson, S. Debei, M. De Cecco, S. Fornasier, M. Fulle, C. Güttler, W.-H. Ip, J. R. Kramm, M. Küppers, L. M. Lara, M. Lazzarin, J. J. Lopez Moreno, F. Marzari, G. Naletto, N. Oklay, N. Thomas, C. Tubiana, and K.-P. Wenzel. The global shape, density and rotation of comet 67p/churyumov-gerasimenko from preperihelion rosetta/osiris observations. *Icarus*, 277:257–278, 2016.
- [24] S. C. Lowry, P. R. Weissman, S. R. Duddy, B. Rozitis, A. Fitzsimmons, S. F. Green, M. D. Hicks, C. Snodgrass, S. D. Wolters, and S. R. Chesley. The internal structure of asteroid (25143) itokawa as revealed by detection of YORP spin-up. *Astronomy and Astrophysics*, 562:A48, 2014.
- [25] JW McMahon, DJ Scheeres, SG Hesar, D Farnocchia, S Chesley, and D Lauretta. The osiris-rex radio science experiment at bennu. *Space Science Reviews*, 214(1):1–41, 2018.
- [26] SR Chesley, AS French, AB Davis, RA Jacobson, M Brozović, D Farnocchia, S Selznick, AJ Liounis, CW Hergenrother, MC Moreau, et al. Trajectory estimation for particles observed in the vicinity of (101955) bennu. *Journal of Geophysical Research: Planets*, 125(9):e2019JE006363, 2020.
- [27] D. S. Lauretta and OSIRIS-REx team. An overview of the OSIRIS-REx asteroid sample return mission. In *Lunar and Planetary Science Conference*, 2012.

- [28] Scott Evans, William Taber, Theodore Drain, Jonathon Smith, Hsi-Cheng Wu, Michelle Guevara, Richard Sunseri, and James Evans. Monte: The next generation of mission design and navigation software. *CEAS Space Journal*, 10(1):79–86, 2018.
- [29] OS Barnouin, MG Daly, EE Palmer, RW Gaskell, JR Weirich, CL Johnson, MM Al Asad, JH Roberts, ME Perry, HCM Susorney, et al. Shape of (101955) bennu indicative of a rubble pile with internal stiffness. *Nature geoscience*, 12(4):247–252, 2019.
- [30] Yu Takahashi and DJ Scheeres. Small body surface gravity fields via spherical harmonic expansions. *Celestial Mechanics and Dynamical Astronomy*, 119(2):169–206, 2014.
- [31] Mary Knapp, Sara Seager, Brice-Olivier Demory, Akshata Krishnamurthy, Matthew W Smith, Christopher M Pong, Vanessa P Bailey, Amanda Donner, Peter Di Pasquale, Brian Campuzano, et al. Demonstrating high-precision photometry with a cubesat: Asteria observations of 55 cancri e. *The Astronomical Journal*, 160(1):23, 2020.
- [32] William M Owen Jr. Methods of optical navigation. 2011.
- [33] MG Daly, OS Barnouin, C Dickinson, J Seabrook, CL Johnson, G Cunningham, T Haltigin, D Gaudreau, C Brunet, I Aslam, et al. The osiris-rex laser altimeter (ola) investigation and instrument. *Space Science Reviews*, 212(1):899–924, 2017.
- [34] John Southworth, TC Hinse, MJ Burgdorf, M Dominik, Allan Hornstrup, UG Jørgensen, C Liebig, Davide Ricci, CC Thöne, T Anguita, et al. High-precision photometry by telescope defocussing–ii. the transiting planetary system wasp-4. *Monthly Notices of the Royal Astronomical Society*, 399(1):287–294, 2009.
- [35] Andrew J Liounis, Jeffrey L Small, Jason C Swenson, Joshua R Lyzhoft, Benjamin W Ashman, Kenneth M Getzandanner, Michael C Moreau, Coralie D Adam, Jason M Leonard, Derek S Nelson, et al. Autonomous detection of particles and tracks in optical images. *Earth and Space Science*, 7(8):e2019EA000843, 2020.
- [36] B. Hockman, A. Frick, I. A. D. Nesnas, and M. Pavone. Design, control, and experimentation of internally-actuated rovers for the exploration of low-gravity planetary bodies. *Journal of Field Robotics*, 34(1):5–24, 2016.
- [37] B. H. Wilcox and R. M. Jones. The MUSES-CN nanorover mission and related technology. In *IEEE Aerospace Conference*, 2000.
- [38] A. Parness, N. Abcouwer, C. Fuller, N. Wiltsie, J. Nash, and B. Kennedy. LEMUR 3: A limbed climbing robot for extreme terrain mobility in space. In *Proc. IEEE Conf. on Robotics and Automation*, 2017.
- [39] C. Rossi, P. M. Cunio, F. Alibay, J. Morrow, S. L. Nothnagel, T. Steiner, C. J. Han, E. Lanford, and J. A. Hoffman. TALARIS project update: Overview of flight testing and development of a prototype planetary surface exploration hopper. *Acta Astronautica*, 81(1):348–357, 2012.
- [40] B. C. Clark, J. M. Sunshine, M. F. A’Hearn, A. L. Cochran, T. L. Farnham, W. M. Harris, T.J. McCoy, and J. Veverka. Comet hopper: A mission concept for exploring the heterogeneity of comets. In *Asteroids, Comets, Meteors*, 2008.
- [41] A. F. J. Abercromby, M. L. Gernhardt, S. P. Chappell, D. E. Lee, and A. S. Howe. Human exploration of phobos. In *IEEE Aerospace Conference*, 2015.

- [42] R. Z. Sagdeev and A. V. Zakharov. Brief history of the phobos mission. *Nature*, 341(6243):581–585, 1989.
- [43] K. Yoshida. The jumping tortoise: A robot design for locomotion on micro-gravity surface. In *Int. Symp. on Artificial Intelligence, Robotics and Automation in Space*, 1999.
- [44] P. Fiorini and J. Burdick. The development of hopping capabilities for small robots. *Autonomous Robots*, 14(2):239–254, 2003.
- [45] K. Schindler, C. A. Thomas, V. Reddy, A. Weber, S. Gruska, and S. Fasoulas. PANIC – a surface science package for the in situ characterization of a near-earth asteroid. *Acta Astronautica*, 68(11-12):1800–1810, 2011.
- [46] T. Yoshimitsu, T. Kubota, I. Nakatani, T. Adachi, and H. Saito. Micro-hopping robot for asteroid exploration. *Acta Astronautica*, 52(2–6):441–446, 2003.
- [47] C. Dietze, F. Herrmann, S. Kuß, C. Lange, M. Scharringhausen, L. Witte, T. van Zoest, and H. Yano. Landing and mobility concept for the small asteroid lander MASCOT on asteroid 1999 JU3. In *Int. Astronautical Congress*, 2010.
- [48] E. Adams, E. Hohlfeld, S. Hill, B. Wilhelm, A. Sharma, B. Williams, and S. Liang. Asteroid in situ exploration using planetary object geophysical observer (POGO). In *IEEE Aerospace Conference*, 2016.
- [49] S. Shimoda, T. Kubota, and I. Nakatani. New mobility system based on elstatic energy under microgravity. In *Proc. IEEE Conf. on Robotics and Automation*, 2002.
- [50] K.-H. Glassmeier, H. Boehnhardt, E. Koschny, D. and Kührt, and I. Richter. The rosetta mission: flying towards the origin of the solar system. *Space Science Reviews*, 128(1):1–21, 2007.
- [51] E. Hand. Philae probe makes bumpy touchdown on a comet. *Science*, 346(6212):900–901, 2014.
- [52] Alfred Anthony Rizzi, Michael Patrick Murphy, John Joseph Giarratana, Matthew David Malchano, Christian Allen Weagle, and Chris Aaron Richburg. Hopping robot, September 30 2014. US Patent 8,849,451.
- [53] B. Hockman, R. G. Reid, I. A. D. Nesnas, and M. Pavone. Experimental methods for mobility and surface operations of microgravity robots. In *Int. Symp. on Experimental Robotics*, 2016.
- [54] Tetsuo Yoshimitsu, Jun’ichiro Kawaguchi, Tatsuaki Hashimoto, Takashi Kubota, Masashi Uo, Hideo Morita, and Kenichi Shirakawa. Hayabusa-final autonomous descent and landing based on target marker tracking. *Acta Astronautica*, 65(5-6):657–665, 2009.
- [55] Dieter Mehrholz, L Leushacke, W Flury, R Jehn, H Klinkrad, and M Landgraf. Detecting, tracking and imaging space debris. *ESA Bulletin(0376-4265)*, (109):128–134, 2002.
- [56] J. A. Atchison and R. H. Mitch. Asteroid flyby gravimetry via target tracking. In *AIAA/AAS Space Flight Mechanics Meeting*, 2015.
- [57] Kohei Fujimoto, Nathan Stacey, and James M Turner. Stereoscopic image velocimetry as a measurement type for autonomous asteroid gravimetry. In *AIAA/AAS Astrodynamics Specialist Conference*, page 5566, 2016.

- [58] Bruce Hapke. Bidirectional reflectance spectroscopy 7: The single particle phase function hockey stick relation. *Icarus*, 221(2):1079–1083, 2012.
- [59] Takushi Tanaka, Yoshiyuki Kawamura, and Takakazu Tanaka. Development and operations of nano-satellite fitsat-1 (niwaka). *Acta Astronautica*, 107:112–129, 2015.
- [60] Christopher M Pong and Matthew W Smith. Camera modeling, centroiding performance, and geometric camera calibration on asteria. In *2019 IEEE Aerospace Conference*, pages 1–17. IEEE, 2019.